

Active Intent Disambiguation for Shared Control Robots

Deepak E. Gopinath^{ID} and Brenna D. Argall

Abstract—Assistive shared-control robots have the potential to transform the lives of millions of people afflicted with severe motor impairments. The usefulness of shared-control robots typically relies on the underlying autonomy’s ability to infer the user’s needs and intentions, and the ability to do so *unambiguously* is often a limiting factor for providing appropriate assistance confidently and accurately. The contributions of this paper are four-fold. First, we introduce the idea of intent disambiguation via control mode selection, and present a mathematical formalism for the same. Second, we develop a control mode selection algorithm which selects the control mode in which the user-initiated motion helps the autonomy to *maximally disambiguate* user intent. Third, we present a pilot study with eight subjects to evaluate the efficacy of the disambiguation algorithm. Our results suggest that the disambiguation system (a) helps to significantly reduce task effort, as measured by number of button presses, and (b) is of greater utility for more limited control interfaces and more complex tasks. We also observe that (c) subjects demonstrated a wide range of disambiguation request behaviors, with the common thread of concentrating requests early in the execution. As our last contribution, we introduce a novel field-theoretic approach to *intent inference* inspired by dynamic field theory that works in tandem with the disambiguation scheme.

Index Terms—Assistive robotics, shared autonomy, intent inference, intent disambiguation.

I. INTRODUCTION

ASSISTIVE and rehabilitation machines—such as robotic arms and smart wheelchairs—have the potential to transform the lives of millions of people with severe motor impairments [1]. With rapid technological advancements in the domain of robotics these machines have become more capable and complex, and with this complexity the control of these machines has become a greater challenge.

The standard usage of these assistive machines relies on manual teleoperation typically enacted through a control interface such as a joystick. However, the greater the motor

impairment of the user, the more limited are the interfaces available for them to use. These interfaces (for example, sip-and-puffs and switch-based head arrays) are low-dimensional, at times discrete, and can typically only operate in subsets of the entire control space (referred to as *control modes*). The dimensionality mismatch between the interface and the robot’s controllable degrees-of-freedom (DoF) necessitates the user to switch between control modes during teleoperation and has been shown to add to the cognitive and physical burden of operation and to affect task performance negatively [2].

The introduction of robotics *autonomy* to these assistive machines can alleviate some of the above-mentioned issues. More specifically, with *shared* autonomy the task responsibility is shared between the user and the underlying autonomy. However, for autonomy to be effective, it needs to have a good understanding of the user’s needs and intentions. That is, *intent inference* is critical to ensure appropriate assistance.

In this work, we consider use-case scenarios in which the autonomy’s inference of user intent is exclusively informed by the human’s control commands issued via the control interface. As an example, in the domain of assistive robotic manipulation, these control commands are typically mapped to the end-effector (or joint) velocities and result in robot motion. Motion carries information regarding underlying intent. However, intent inference becomes particularly challenging when the user input is low-dimensional and sparse—as is the case with the more limited interfaces available to those with severe motor impairments—because the robot motion will likely be more discontinuous and jagged, and carries less direct information regarding the underlying human intent. While to augment the human-robot system with high-fidelity sensors could enhance the autonomy’s intent inference capabilities, for reasons of user adoption and cost, within the assistive domain we intentionally design our assistance add-ons to be as invisible and close to the manual system as possible. The need for intent *disambiguation* arises as the autonomy needs to reason about all possible goals before issuing appropriate assistance commands.

Our key insight in this work is that user control commands issued in certain control modes are *more intent expressive* than others and therefore may help the autonomy to improve inference accuracy. More specifically, in this work we investigate how the selection of a subset of the operational control dimensions or modes improves the intent inference and disambiguation capabilities of the robot. The idea is that earlier and more accurate inference enables the autonomy to assist the human earlier and more effectively, and thereby

Manuscript received December 2, 2018; revised April 21, 2019, July 29, 2019, and November 16, 2019; accepted December 18, 2019. Date of publication April 14, 2020; date of current version June 5, 2020. This work was supported by the National Science Foundation under Grant CNS 1544741. (Corresponding author: Deepak E. Gopinath.)

Deepak E. Gopinath is with the Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208 USA, and also with the Shirley Ryan AbilityLab, Chicago, IL 60611 USA (e-mail: deepakedakkattilgopinath2015@u.northwestern.edu).

Brenna D. Argall is with the Department of Mechanical Engineering, Computer Science and Physical Medicine and Rehabilitation, Northwestern University, Evanston, IL 60208 USA, and also with the Shirley Ryan AbilityLab, Chicago, IL 60611 USA.

Digital Object Identifier 10.1109/TNSRE.2020.2987878

improve overall task performance. This is important especially in the domain of assistive robotics, wherein the purpose of the autonomy is to bridge gaps in control proficiency that result from human impairments and limited control interfaces. The main contributions of this work are as follows:

- 1) We introduce the idea of intent disambiguation via control mode selection, and present a mathematical formalism that frames disambiguation as the problem of selecting the control mode able to maximally disambiguate between goals.
- 2) We develop a control mode selection algorithm which selects the control mode *for* the user, in which the user-initiated motion will help the autonomy to *maximally disambiguate* intent by eliciting more *intent expressive* control commands from the human.
- 3) We present results from a pilot study conducted to evaluate the efficacy of the disambiguation algorithm.
- 4) We propose a novel field-theoretic approach to intent inference based on ideas from *dynamic field theory* in which the time evolution of the probability distribution over goals is specified as a continuous-time constrained dynamical system that obeys the principle of maximum entropy in the absence of user control commands.

In Section II we present an overview of relevant research in the areas of shared autonomy in assistive robotics, intent inference, and synergies in human-robot interaction. Section III presents our mathematical formalism for intent disambiguation. The disambiguation and intent inference algorithms are outlined in Sections IV and V respectively. The study design and experimental methods are discussed in Section VI followed by results in Section VII. Discussion and conclusions are presented respectively in Sections VIII and IX.

II. RELATED WORK

This section provides an overview of related research in the domains of shared autonomy in assistive robotics, intent inference in human-robot interaction, and synergies in human-robot interaction.

Shared control in assistive systems aims to reduce the user's cognitive and physical burden during task execution, typically without having the user relinquish complete control [3]–[6]. In order to offset the drop in task performance due to shifting focus from the task at hand to switching between different control modes, various mode switch assistance paradigms have been proposed. A simple time-optimal mode switching scheme has shown to improve task performance [2], [7].

Shared-control systems often require a good estimate of the human's intent—for example, their intended reaching target in a manipulation task or a target goal location in a navigation task [8]. Intent can be explicitly communicated by the user [9] via various modalities such as laser pointers, click interfaces, and in some cases natural language [10]. Intent can also be inferred from the user's control signals and other environmental cues using various algorithms [11]. Within the context of shared autonomy a Bayesian scheme for user intent prediction models the user within the Markov Decision Process framework [12]–[14] and is typically assumed to

be noisily optimizing some cost function for their intended goal. In low-dimensional spaces, this cost function can be learned from expert demonstrations using Inverse Reinforcement Learning [15].

For high-dimensional spaces, such as that of robotic manipulation, learning cost functions that generalize well over the entire space requires large number of samples. In such cases, heuristic cost functions, such as sum of squared velocities along a trajectory, have been found to be useful for goal prediction [16]. Simple heuristic approaches can also be used to find direct mappings from instantaneous cues and the underlying human intention. Heuristic approaches can incorporate domain-specific knowledge easily and are computationally inexpensive, though the trade-off for this simplicity is not being sophisticated enough to incorporate histories of states and actions, making them less robust to external noise. Instantaneous confidence functions for estimating the intended reaching target are employed with success on multiple robotic manipulation systems [5], [17]. In our work we develop an inference algorithm that updates the belief over goals using ideas from dynamic field theory in which the histories of states and actions are incorporated using a single time-scale parameter and robustness to noise is ensured via recurrent self-interactions that stabilize the dynamical system.

From the robot's perspective, the core idea behind our intent disambiguation approach is one of “*Help Me, Help You*”—that is, if the user can help the autonomy with disambiguation via more intent-expressive actions, then the autonomy in turn can provide appropriate task assistance more swiftly and accurately. In human-robot interaction, the legibility and predictability of robot motion *to* the human is investigated [18] with various techniques to generate legible robot motion proposed [19]. Our work instead investigates the idea of *inverse legibility* [20] in which the assistance scheme is intended to bring out more legible intent-expressive control commands *from* the human.

III. MATHEMATICAL FORMALISM FOR INTENT DISAMBIGUATION

This section frames intent disambiguation as a problem of determining the control mode able to maximally disambiguate between goals.

A. Notation

Let \mathcal{G} be the set of all candidate goals with $n_g = |\mathcal{G}|$ and g^i refer to the i^{th} goal with $i \in [1, 2, \dots, n_g]$. A *goal* in this context represents the human's underlying intent. Specifically, in assistive robotic manipulation, as the robotic arm first must reach toward and grasp discrete objects in the environment, intent inference is the estimation of the belief over all possible discrete goals (objects) in the environment. At any time t , the autonomy maintains a probability distribution over goals denoted by $\mathbf{p}(t)$ such that $\mathbf{p}(t) = [p^1(t), p^2(t), \dots, p^{n_g}(t)]^T$ where $p^i(t)$ denotes the probability associated with goal g^i . The probability $p^i(t)$ represents the robot's belief that goal g^i is the human's intended goal.

Let \mathcal{K} be the set of all controllable dimensions of the robot and k^i represent the i^{th} control dimension where $i \in [1, 2, \dots, n_k]$ with $n_k = |\mathcal{K}|$. The limitations of the control interface necessitate \mathcal{K} to be partitioned into control modes. Let \mathcal{M} be the set of all control modes with $n_m = |\mathcal{M}|$. Additionally, let m^i refer to the i^{th} control mode where $i \in [1, 2, \dots, n_m]$. Each control mode m^i is a subset of \mathcal{K} such that $\bigcup_{i=1}^{n_m} m^i = \mathcal{K}$.¹

In this work, we assume a kinematic model for the robot and the kinematic state (the robot's end-effector pose) at any time t is denoted as $\mathbf{x}_r(t) \in \mathbb{R}^3 \times \mathbb{S}^3$ and consists of a position and orientation component, where \mathbb{S}^3 is the space of all unit quaternions. The pose for goal $g \in \mathcal{G}$ is denoted as $\mathbf{x}_g \in \mathbb{R}^3 \times \mathbb{S}^3$. The control command issued by the human via the control interface is denoted as \mathbf{u}_h and is mapped to the Cartesian velocity of the robot's end-effector. For a 6-DoF robotic arm, $\mathbf{u}_h \in \mathbb{R}^6$. The autonomous control policy generates an autonomy control command which is denoted as $\mathbf{u}_a \in \mathbb{R}^6$. The control command issued to the robot, which is a synthesis of \mathbf{u}_h and \mathbf{u}_a is denoted as $\mathbf{u} \in \mathbb{R}^6$. The control command that corresponds to a unit velocity vector along control dimension k is denoted as \mathbf{e}^k .

B. Disambiguation Metric

The disambiguation metrics that we develop are *heuristic* measures that characterize the intent disambiguation capabilities of robot motion along particular control dimensions within particular control modes. Specifically, we define disambiguation metric $D_k \in \mathbb{R} \forall k \in \mathcal{K}$. We further explicitly denote disambiguation measures for both positive and negative motions along k as D_k^+ and D_k^- respectively. We also define a disambiguation metric $D_m \in \mathbb{R}$ for each control mode $m \in \mathcal{M}$. By virtue of design, the disambiguation metric D_m is a measure of how useful robot motion would be to the autonomy's ability to perform intent inference if the user were to operate the robot in control mode m . In this work our computation of D_k depends on four features (denoted as Γ_k , Ω_k , Λ_k and Υ_k), that capture different aspects of the *shape* of a projection of the probability distribution over intent. These projections and computations are described in detail in Section IV-A and Section IV-B, and as pseudocode in Algorithm 1.

IV. DISAMBIGUATION ALGORITHM

This section describes the computation of D_k and D_m contributed in this article, as well as our intent disambiguation algorithm.

A. Forward Projection of $\mathbf{p}(t)$

The first step in the computation of D_k is a model-based forward projection of the probability distribution $\mathbf{p}(t)$ from the current time t_a to times t_b and t_c (Algorithm 1, line 4)

¹Note that a dimension $k \in \mathcal{K}$ can be an element of multiple control modes.

Algorithm 1 Intent Disambiguation

Require: $\mathbf{p}(t_a), \mathbf{x}_r(t_a), \Delta t, t_a < t_b < t_c, \Theta$

- 1: **for** $k = 0 \dots n_k$ **do**
- 2: Initialize $D_k = 0, t = t_a, \mathbf{u}_h = \mathbf{e}^k$
- 3: **while** $t \leq t_c$ **do**
- 4: $\mathbf{p}_k(t + \Delta t) \leftarrow \text{UpdateIntent}(\mathbf{p}_k(t), \mathbf{u}_h; \Theta)$
- 5: $\mathbf{x}_r(t + \Delta t) \leftarrow \text{SimulateKinematics}(\mathbf{x}_r(t), \mathbf{u}_h)$
- 6: **if** $t = t_b$ **then**
- 7: Compute $\Gamma_k, \Omega_k, \Lambda_k$
- 8: **end if**
- 9: **if** $t = t_c$ **then**
- 10: Compute Υ_k
- 11: **end if**
- 12: $t \leftarrow t + \Delta t$
- 13: **end while**
- 14: Compute D_k
- 15: **end for**

where $t_a < t_b < t_c$.² We consider two future times in order to compute short-term (t_b) and long-term (t_c) evolutions of the probability distribution. The application of unit velocity \mathbf{e}^k results in probability distributions $\mathbf{p}_k^+(t_b)$ and $\mathbf{p}_k^+(t_c)$, and $-\mathbf{e}^k$ results in $\mathbf{p}_k^-(t_b)$ and $\mathbf{p}_k^-(t_c)$, where the subscript k captures the fact that the projection is the result of the application of a control command only along control dimension k .

B. Features of D_k

To compute our control dimension disambiguation metric, we design four features that encode different aspects of the *shape* of the probability distribution as it evolves under motion in a specific control dimension k . For each control dimension k , each of the four features is computed for projections along both positive and negative directions independently. The four features are computed in lines 7 and 10 in Algorithm 1.

1) *Maximum*: The maximum of the projected probability distribution $\mathbf{p}_k(t_b)$ is a good measure of the robot's *overall certainty* in accurately predicting human intent. We define the distribution maximum as

$$\Gamma_k = \max_{1 \leq i \leq n_g} p_k^i(t_b) \quad (1)$$

(i.e., the statistical mode of this discrete distribution). A higher value implies that the robot has a greater confidence in its prediction of the human's intended goal.

2) *Pairwise separation*: More generally, disambiguation accuracy benefits from a larger separation, Λ_k , between goal probabilities. The quantity Λ_k is computed as the *sum of the pairwise distances* between the n_g probabilities.

$$\Lambda_k = \sum_{i=1}^{n_g} \sum_{j=i}^{n_g} |p_k^i(t_b) - p_k^j(t_b)| \quad (2)$$

²*UpdateIntent()* in Line 4 is implemented using Equation 6 discussed in detail in Section V-B. *SimulateKinematics()* assumes that the end-effector kinematics is same as that of a point-like robot. All parameters which affect the computation of $\mathbf{p}(t)$ are denoted as Θ .

Λ_k is particularly helpful if the difference between the largest probabilities fails to disambiguate.

3) *Difference between maxima*: Disambiguation accuracy benefits from greater differences between the first and second most probable goals. This difference is denoted as

$$\Omega_k = \max(\mathbf{p}_k(t_b)) - \max(\mathbf{p}_k(t_b) \setminus \max(\mathbf{p}_k(t_b))) \quad (3)$$

and Ω_k becomes particularly important when the distribution has multiple modes and a single measure of maximal certainty (Γ_k) alone is not sufficient for successful disambiguation.

4) *Gradients*: Γ_k , Ω_k and Λ_k are local measures that encode shape characteristics of the short-term temporal projections of the probability distribution over goals. However, the quantity $\mathbf{p}_k(t)$ can undergo significant changes upon long-term continuation of motion along control dimension k . The spatial gradient of $\mathbf{p}_k(t)$ encodes this propensity for change and is approximated by

$$\frac{\partial \mathbf{p}_k(t)}{\partial x_k} \simeq \frac{\mathbf{p}_k(t_c) - \mathbf{p}_k(t_b)}{x_k(t_c) - x_k(t_b)}$$

where x_k is the component of robot's projected displacement along control dimension k . The greater the difference between individual spatial gradients, the greater will the probabilities deviate from each other, thereby helping in disambiguation. In order to quantify the "spread" of gradients we define Υ_k as

$$\Upsilon_k = \sum_{i=1}^{n_g} \sum_{j=i}^{n_g} \left| \frac{\partial p_k^i(t)}{\partial x_k} - \frac{\partial p_k^j(t)}{\partial x_k} \right| \quad (4)$$

where $|\cdot|$ denotes the absolute value.

5) *Computation of D_k and D_m* : The individual features Γ_k , Ω_k , Λ_k and Υ_k are combined to compute D_k in such a way that, by design, higher values of D_k imply greater disambiguation capability for the control dimension k . More specifically,

$$D_k = \underbrace{w \cdot (\Gamma_k \cdot \Lambda_k \cdot \Omega_k)}_{\text{short-term}} + \underbrace{(1-w) \cdot \Upsilon_k}_{\text{long-term}} \quad (5)$$

where w is a task-specific weight that balances the contributions of the short-term and long-term components. In our implementation, w is empirically set to 0.5. Equation 5 is computed twice, once in each of the positive (\mathbf{e}^k) and negative directions ($-\mathbf{e}^k$) along k , and the results (D_k^+ and D_k^-) are then summed to compute D_k .

The computation of D_k is performed for each control dimension $k \in \mathcal{K}$. The disambiguation metric D_m for control mode m then is calculated as

$$D_m = \sum_{k \in m} D_k$$

and the control mode with highest disambiguation capability m^* is given by $m^* = \operatorname{argmax}_m D_m$ while $k^* = \operatorname{argmax}_k D_k$ gives the control dimension with highest disambiguation capability k^* . Disambiguation mode m^* is the mode that the algorithm chooses for the human to better estimate their intent.

V. INTENT INFERENCE

Since the disambiguation power of our algorithm is closely linked to the fidelity of the underlying intent inference mechanism, in this section, we propose a novel intent inference scheme inspired by *dynamic field theory*. By having the autonomy maintain a probability distribution over goals, we implicitly model the human as a Partially Observable Markov Decision Process (POMDP) in which all the uncertainty in the user's state is concentrated in the user's intended goal. By maintaining and updating a probability distribution over goals the autonomy can reason about the human's latent state (internal goal) during trial execution. Inference over goal states typically is done using a recursive Bayesian belief update which determines how the distribution evolves over time. Here we introduce a novel approach to compute the time evolution of a probability distribution over goals as a dynamical system with constraints which serves as an alternative to the recursive Bayesian update scheme.

A. Dynamic Field Theory

In Dynamic Field Theory (DFT) [21], variables of interest are treated as dynamical state variables. To represent the information about these variables requires two dimensions: one which specifies the value the variables can attain and the other which encodes the *activation level* or the amount of information about a particular value. These *activation fields* (also known as dynamic neural fields) are analogous to probability distributions defined over a random variable.

Following Amari's formulation [22] the dynamics of an activation field $\phi(x, t)$ are given by

$$\tau \frac{\partial \phi(x, t)}{\partial t} = \int dx' b(x-x') \sigma(\phi(x', t)) - \phi(x, t) + h + S(x, t)$$

where x denotes the variable of interest, t is time, τ is the time-scale parameter, $b(x-x')$ is the interaction kernel, $\sigma(\phi)$ is a sigmoidal nonlinear threshold function, h is the constant resting level and $S(x, t)$ is the external input. The interaction kernel mediates how activations at all other field sites x' drive the activation level at x . Two types of interactions are possible: excitatory (when interaction is positive) which drives up the activation, and inhibitory (when the interaction is negative) which drives the activation down.

Historically, dynamic neural fields were conceived to explain cortical population neuronal dynamics based on the hypothesis that the excitatory and inhibitory neural interactions between local neuronal pools form the basis of cortical information processing. These activation fields possess unique characteristics that make them ideal candidates for modeling the time evolution of $\mathbf{p}(t)$. First, a peak in the activation field can be *sustained* even in the absence of external input due to the recurrent interaction terms. Second, information from the past can be *preserved* over much larger time scales quite easily by tuning the time-scale parameter thereby endowing the fields with memory. Third, the activation fields are *robust* to disturbance and noise in the external input [23]. We harness these characteristics to specify a model for smooth temporal evolution of $\mathbf{p}(t)$ in the next section.

B. Field-Theoretic Intent Inference

Our insight is to use the framework of dynamic neural fields to specify the time evolution of the probability distribution $\mathbf{p}(t)$, in which we treat the individual goal probabilities $p^i(t)$ as constrained dynamical state variables such that $p^i(t) \in [0, 1]$ and $\sum_1^{n_g} p^i(t) = 1$. We refer to this approach as the *field-theoretic intent inference*.

The full specification of the field is given by

$$\begin{aligned} \frac{\partial \mathbf{p}(t)}{\partial t} = \frac{1}{\tau} & \left[\underbrace{-\mathbf{P}_{n_g \times n_g}^T \cdot \mathbf{p}(t)}_{\text{goal transition dynamics}} + \underbrace{\frac{1}{n_g} \cdot \mathbb{1}_{n_g}}_{\text{rest state}} \right] \\ & + \underbrace{\lambda_{n_g \times n_g} \cdot \sigma(\boldsymbol{\xi}(\mathbf{u}_h; \boldsymbol{\Theta}))}_{\text{excitatory + inhibitory}} \end{aligned} \quad (6)$$

where time-scale parameter τ determines the memory capacity and decay behavior, $\mathbf{P}_{n_g \times n_g}$ is the state transition matrix for the embedded Markov chain that models the goal transitions as jump processes, $\mathbb{1}_{n_g}$ is a vector of dimension n_g containing all ones, \mathbf{u}_h is the human control input and $\boldsymbol{\Theta}$ represents all other task-relevant features, λ is the control matrix that controls the excitatory and inhibitory aspects, $\boldsymbol{\xi}$ is a function that encodes the nonlinearity through which human control commands and task features affect the time evolution, and σ is a biased sigmoidal nonlinearity given by $\sigma(\boldsymbol{\xi}) = \frac{1}{1+e^{-\boldsymbol{\xi}}} - 0.5$.

Our design of $\boldsymbol{\xi}$ is informed by what features of the human control input and environment effectively capture the human's underlying intent. We choose the *directedness* of the robot motion towards a goal, the *agreement* between the human and robot autonomy commands, and the *proximity* to a goal. The *directedness* component looks at the shortest straight line path towards a goal g , whereas the *agreement* serves as an indicator of how similar (measured as a dot product) the human and autonomy signals are to each other. One dimension i of $\boldsymbol{\xi}$ is defined as

$$\begin{aligned} \xi^i(\mathbf{u}_h; \boldsymbol{\Theta}) = & \underbrace{\frac{1+\eta}{2}}_{\text{directedness}} + \underbrace{\mathbf{u}_h^{\text{rot}} \cdot \mathbf{u}_{a,g^i}^{\text{rot}}}_{\text{agreement}} \\ & + \underbrace{\max\left(0, 1 - \frac{\|\mathbf{x}_{g^i} - \mathbf{x}_r\|}{R}\right)}_{\text{proximity}} \end{aligned}$$

where *trans* and *rot* refer to the translational and rotational components of a command \mathbf{u} or position \mathbf{x} , $\eta = \frac{\mathbf{u}_h^{\text{trans}} \cdot (\mathbf{x}_{g^i} - \mathbf{x}_r)^{\text{trans}}}{\|\mathbf{u}_h^{\text{trans}}\| \|\mathbf{x}_{g^i} - \mathbf{x}_r\|}$, \mathbf{u}_{a,g^i} is the robot autonomy command for reaching goal g^i , R is the radius of the sphere beyond which the proximity component is always zero, \cdot is the Euclidean norm and $\boldsymbol{\Theta} = \{\mathbf{x}_r, \mathbf{x}_{g^i}, \mathbf{u}_{a,g^i}\}$. At every time-step, constraints on $p^i(t)$ are enforced such that $\mathbf{p}(t)$ is a valid probability distribution. The most confident goal g^* is computed as $g^* = \text{argmax}_i p^i(t) \forall i \in [1, \dots, n_g]$.

C. Field-Theoretic Intent Inference For Assistive Robotics

In our work, the autonomy's inference of user intent solely relies on user control commands. In the domain of assistive

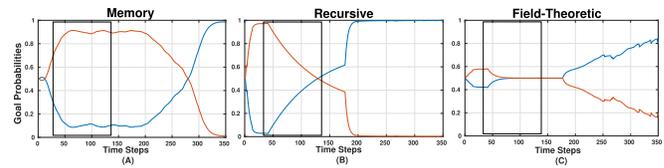


Fig. 1. Inference comparison: Goal probabilities (blue and orange lines) for (A) Memory-Based Prediction, (B) Recursive Bayesian Belief Update and (C) Field-Theoretic Inference in a scene with two goals. Black rectangular boxes indicate times of zero control velocity, with varying effects on the inference schemes: (A) little change (since the cost function is purely distance-based), and convergence to (B) the stationary distribution of the goal transition matrix \mathbf{P} and (C) the uniform distribution as dictated by the principle of maximum entropy.

robotics, it is quite often the case that the user input is highly discontinuous (due to fatigue, motor impairments, stoppage for mode switches, *et cetera*). Therefore, it is important to reason about belief over goals also in the *absence* of useful information.

According to the *principle of maximum entropy*, in the absence of testable information (no control commands issued by the user and a uniform global prior), the belief should converge to a uniform distribution. In the absence of \mathbf{u}_h , using Equation 6 and an appropriately chosen time-scale parameter τ , $\mathbf{p}(t)$ converges to a uniform distribution by correctly ignoring outdated information. The rate at which the distribution decays to a uniform distribution is controlled by τ .

By contrast, the standard discrete-time recursive belief update equation as implemented in [11] is

$$p(g^t | \mathbf{u}_h^t) = \eta p(\mathbf{u}_h^t | g^t) \sum_{g^{t-1} \in \mathcal{G}} p(g^t | g^{t-1}) p(g^{t-1} | \mathbf{u}_h^{t-1})$$

where η is a normalization factor, $p(\mathbf{u}_h | g)$ is a likelihood function, and $p(g^t | g^{t-1})$ is the goal transition probability. In the recursive belief update, when $\mathbf{u}_h = 0$ and the likelihood function is uniform, it can be shown that the posterior distribution over goals converges to the stationary distribution of the goal transition matrix. The stationary distribution is not necessarily uniform and can introduce unwanted biases in the inference.

Knowledge of task-level semantics can provide informative global priors that can further improve the accuracy of the inference mechanism. Our field-theoretic approach additionally can encode a task-level global prior in the ‘rest state’ term. For example, in a pick-and-place task, the initial goal distribution could be biased towards the object that needs to be picked.

In order to evaluate the performance of our field-theoretic inference approach a quantitative comparison to (1) memory-based prediction [16] and (2) recursive belief updating [11] was implemented using point robot simulation in \mathbb{R}^3 . The human was modeled as issuing a control command that noisily optimizes a straight-line path towards the intended goal. Signal dropout was simulated by randomly zeroing out control commands and τ was set to be 10. Additionally, \mathbf{u}_h was set to be zero for a randomly chosen section of each trial in order to compare the convergence behavior of different approaches. The number of goals varied between three and five. Goal transitions were randomly sampled every five to eight time steps. The average length of the simulated trajectories was



Fig. 2. A 2-axis joystick (left) and switch-based head array (center) and their operational paradigms (right). v and ω indicate the translational and rotational velocities of the end-effector, respectively.

615 time steps. 500 trials were simulated. Inference accuracy was computed as the fraction of total trial time (excluding when $\mathbf{u}_h = 0$) for which an algorithm correctly inferred the ground truth.

Results for field-theoretic inference outperformed memory-based prediction significantly (87.46% vs. 59.15% respectively) and were comparable to recursive belief updating (87.43%). Figure 1 shows an illustrative example of goal inference using the various methods. One can see that when there is no control command issued, the field-theoretic approach alone converges to the maximum entropy uniform distribution.

VI. STUDY METHODS

In this section, we describe the study methods used to evaluate the efficacy of the disambiguation system.

A. Participants

For this study eight subjects were recruited (mean age: 31 ± 11 , 3 males and 5 females). All participants gave their informed, signed consent to participate in the experiment, which was approved by Northwestern University's Institutional Review Board.

B. Hardware

The experiments were performed using the MICO 6-DoF robotic arm (Kinova Robotics, Canada), specifically designed for assistive purposes. The software system was implemented using the Robot Operating System (ROS) and data analysis was performed in MATLAB. The subjects teleoperated the robot using two different control interfaces: a 2-axis joystick and a switch-based head array, controlling the 6D Cartesian velocity of the end-effector (Figure 2). An external button was provided to request the mode switch assistance.

In detail, the joystick generated 2D continuous control signals. Under joystick control the full control space was partitioned into five control modes that were accessed via button presses. The switch-based head array consisted of three switches embedded in a headrest, operated via head movements, and generated 1D discrete signals. Under head array control the full control space was partitioned into seven control modes, the back switch was used to cycle between the different control modes, and the switches to the left and right controlled the motion of the robot's end effector in the positive and negative directions along a selected control dimension.

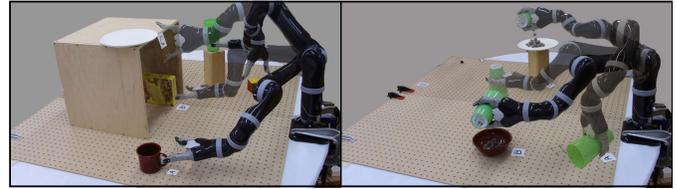


Fig. 3. Study tasks performed by subjects. Left: Single-step reaching task. Right: Multi-step pouring task.

C. Tasks

Two different categories of tasks were evaluated.

1) *Single-Step*: The aim was to reach one of five objects on the table, each with a target orientation (Figure 3, Left).

2) *Multi-Step*: Each trial began with a full cup held by the robot gripper. The task required first that the contents of the cup be poured into one of two containers, and then that the cup be placed at one of the two specified locations and with a particular orientation (Figure 3, Right).

D. Switching Paradigms

Two kinds of mode switching paradigms were evaluated in the study.

1) *Manual*: During task execution the user performed all mode switches.

2) *Disambiguation*: The user either performed a mode switch manually or requested a switch to the *disambiguation* mode. The user was free to issue disambiguation requests at any time during the task execution, upon which the algorithm identified and switched the current control mode to the best disambiguation mode m^* by invoking Algorithm 1. To be clear, the user also was allowed to switch control modes using a manual mode switch at any time as well. The only requirement was that, the user request disambiguation at least once during the task execution.

E. Shared Control

Autonomy assistance was always active for both mode switch assistance paradigms (manual and disambiguation). We used a blending-based shared-control paradigm in which the final robot control command was a linear composition of the human control command and an autonomous control command. With blending the amount of assistance was directly proportional to the probability of the most confident goal g^* , and thus to the strength of the intent inference. The probability distribution over goals, $\mathbf{p}(t)$, was updated using Equation 6 as outlined in Section V-B and the most confident goal was computed as $\text{argmax}_i p^i(t)$. Therefore, if intent inference improved as a result of goal disambiguation, more assistance would be provided by the autonomy.

Specifically, the autonomous control policy generated control command $\mathbf{u}_a \leftarrow f_a(\mathbf{x}_r)$ where $f_a(\cdot) \in \mathcal{F}_a$, and \mathcal{F}_a was the set of all control behaviors corresponding to different tasks. This set could be derived using a variety of techniques such as *Learning from Demonstrations* [24], motion planners [25] or navigation functions [26]. In our implementation, the autonomy's control command was generated using a simple potential field which is defined in all parts of the state

space [27]. Every goal g was associated with a potential field γ_g which treats g as an attractor and all other goals in the scene as repellers. The autonomy command was computed as a summation of the attractor and repeller velocities and operated in the full 6D Cartesian space.

Let $\mathbf{u}_{a,g}$ be the autonomy command associated with goal g . Under blending, the final control command \mathbf{u} issued to the robot then was given by

$$\mathbf{u} = \alpha \cdot \mathbf{u}_{a,g^*} + (1 - \alpha) \cdot \mathbf{u}_h \quad (7)$$

where g^* was the most confident goal. Similar to \mathbf{u}_h , the autonomy command $\mathbf{u}_{a,g^*} \in \mathbb{R}^6$ was mapped to the 6D Cartesian velocity of the end-effector. The blending factor α was a piecewise linear function of the probability $p(g^*)$ associated with g^* and was given by

$$\alpha = \begin{cases} 0 & p(g^*) \leq \rho_1 \\ \frac{\rho_3(p(g^*) - \rho_1)}{\rho_2 - \rho_1} & \text{if } \rho_1 < p(g^*) \leq \rho_2 \\ \rho_3 & \rho_2 < p(g^*) \end{cases}$$

with $\rho_i \in [0, 1] \forall i \in [1, 2, 3]$ and $\rho_2 > \rho_1$. In our implementation, we empirically set $\rho_1 = \frac{1.2}{n_g}$, $\rho_2 = \frac{1.4}{n_g}$ and $\rho_3 = 0.7$.

F. Study Protocol

A within-subjects study was conducted using a fractional factorial design in which the manipulated variables were the tasks, control interfaces, and the switching paradigm conditions. Each subject underwent an initial training period that lasted approximately 30 minutes. The training period consisted of three phases and two different task configurations. The subjects used both interfaces to perform the training tasks.

1) *Phase One*: The subjects were asked to perform a simple reaching motion towards a single goal in the scene. This phase was intended for the subjects to get familiarized with the control interface mappings and teleoperation of the robotic arm.

2) *Phase Two*: Subjects were asked to perform a simple reaching motion towards a single goal in the scene in the presence of blending-based autonomous assistance.

3) *Phase Three*: Subjects were able to explore the disambiguation request feature during a reaching task to observe the effects of the mode switch request and subsequent change in robot assistance. Multiple objects were introduced in the scene. Subjects were explicitly informed that upon a disambiguation request the robot would select a control mode that would help the autonomy determine the subject's intended goal and thereby enable it to assist the user more effectively.

During the testing phase, each subject performed both tasks using both interfaces under the *Manual* and *Disambiguation* paradigms. All trials started in a randomized initial control mode and robot position. The ordering of control interfaces and paradigms was randomized and counterbalanced across all subjects. Three trials were collected for the *Manual* paradigm and five trials for the *Disambiguation* paradigm. On an average, each trial lasted approximately 10-40s depending on the starting position of the robot and the specified reaching target. At the start of each trial, $\mathbf{p}(t)$ was initialized as $\frac{1}{n_g} \cdot \mathbb{1}_{n_g}$. During

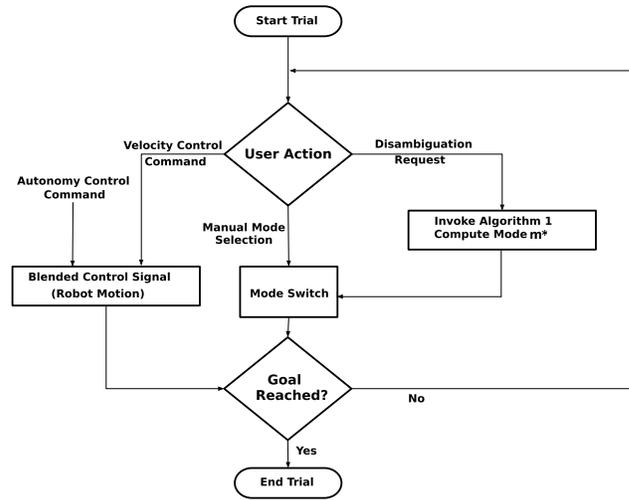


Fig. 4. Flow chart depicting user action sequence during a single trial. The user could issue a (i) velocity control command, resulting in intent inference followed by generation of an autonomy signal and then blended control signal, and causes robot motion or (ii) manual mode switch or (iii) disambiguation request, both resulting in a control mode switch.

the trial as the user teleoperated the robot, $\mathbf{p}(t)$ was updated according to Equation 6 online at each time step. Figure 4 captures how a single trial unfolds in time.

G. Metrics

The objective metrics used for evaluation included the following.

- *Number of mode switches*: The number of times a user switched between various control modes during task execution. This metric captures one of the main factors that contributes to the cognitive and physical effort required for task execution in assistive robotic manipulation [2].
- *Number of disambiguation requests*: The number of times a user pressed the disambiguation request button.
- *Number of button presses*: The sum of *Number of mode switches* and *Number of disambiguation requests*.
- *Skewness*: A higher-order moment used to quantify the asymmetry of any distribution. Used to characterize how much the temporal distribution of disambiguation requests deviates from a uniform distribution.
- *Task completion time*: Time taken to complete the task successfully. This metric is an indicator of how well the human and autonomy work together.

Additionally, at the end of each testing phase, subjective data was gathered via a brief questionnaire. Users were given the following statements regarding the usefulness and capability of the assistance system to rate according to their agreement on a 7-point Likert scale.

- **Q1** - Control modes chosen by the system made task execution easier.
- **Q2** - The robot and I worked together to accomplish the task.
- **Q3** - I liked operating the robot in the control modes chosen by the system.

Subjects were also asked to indicate their preference in the following questions.

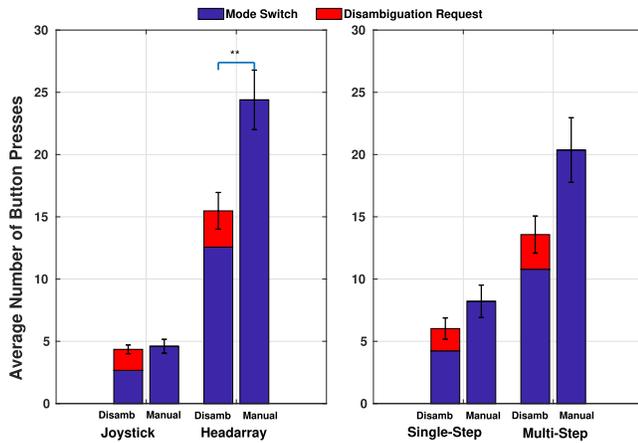


Fig. 5. Average number of button presses, *Disambiguation* and *Manual* paradigms. *Left*: Grouped by control interfaces. *Right*: Grouped by tasks.

- **Q4** - Which interface was the hardest to operate?
- **Q5** - For which interface was the assistance paradigm the most useful?
- **Q6** - Which one of the schemes do you prefer the most?
- **Q7** - Which one of the schemes is the most user-friendly?

VII. RESULTS

In this section we present results from our human-subject study. Our study results indicate that the disambiguation request system is of greater utility for more limited control interfaces and more complex tasks. Subjects demonstrated a wide range of disambiguation request behaviors with a common theme of relying on disambiguation assistance earlier in the trials. Furthermore, the survey results show that operating the robot in the disambiguating mode make task execution easier and that users prefer the *Disambiguation* paradigm to the *Manual* paradigm. Statistical significance is determined using the Wilcoxon Rank-Sum test in where (***) indicates $p < 0.001$, (**) $p < 0.01$ and (*) $p < 0.05$.

A. Impact of Disambiguation

A statistically significant decrease in the number of button presses is observed between the *Manual* and *Disambiguation* paradigms when using the head array (Figure 5, Left). Due to the low-dimensionality of the head array and cyclical nature of mode switching, the number of button presses required for task completion is inherently high. The disambiguation paradigm is helpful in reducing the number of button presses likely due to higher robot assistance that is present in the disambiguating control mode. For the joystick, statistically significant differences between the two paradigms are observed for the number of manual mode switches ($p < 0.05$). However, this gain is offset by the button presses that are required to make disambiguation requests. When grouping by task, the general trend of a decrease in the number of button presses is more pronounced for the more complex multi-step task (Figure 5, Right). Although not statistically significant, we also observe that the autonomy has higher control authority (as measured by α) during the disambiguation trials ($\alpha = 0.27 \pm 0.16$) when

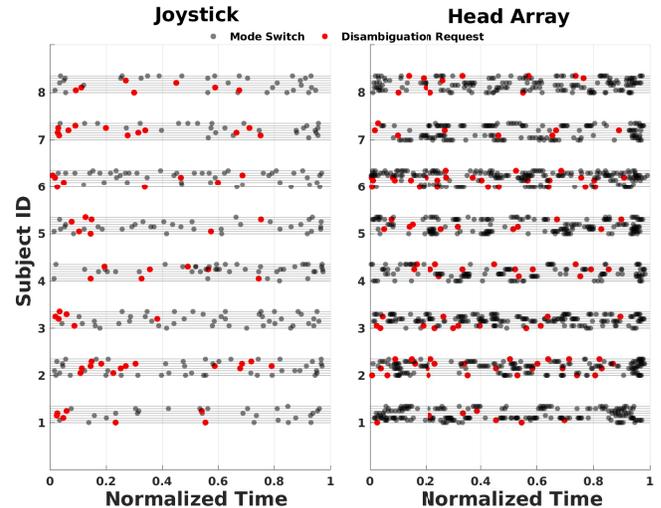


Fig. 6. Temporal pattern of button presses for joystick (left) and head array (Right) during the multi-step task on a trial-by-trial basis for all subjects. For each subject, each light gray horizontal line represents a single trial. Eight trials per subject, for each interface.

compared to the manual trials ($\alpha = 0.25 \pm 0.16$). In Equation 7, larger α allocates more control authority to the autonomy.

These results suggest that disambiguation is more useful as the control interface becomes more limited and the task becomes more complex. Intuitively, intent prediction becomes harder for the robot when the control interface is lower dimensional and does not reveal a great deal about the user's underlying intent. By having the users operate the robot in the disambiguating mode, the robot is able to elicit more intent-expressive control commands from the human which in turn helps in accurate goal inference and subsequently appropriate robot assistance.

B. Temporal Distribution of Disambiguation Requests

In Figure 6 the frequency and density of button presses (disambiguation requests plus mode switches) are much higher for the more limited control interface (head array). We observe that a higher number of disambiguation requests correlates with the more limited interface and complex task. The subjects also demonstrate a diverse range of disambiguation request behaviors, in regards to both (a) when the disambiguation requests are made and (b) with what frequency (e.g., Subject 1 vs. Subject 2, Joystick). The variation between subjects is likely due to different factors such as the user's comfort in operating the robot and understanding of the disambiguating mode's ability to recruit more assistance from the autonomy.

The temporal distribution of disambiguation analyzes *when* the subject requested assistance during the course of a trial. The skewness of the temporal distribution of disambiguation requests reveals a higher concentration of requests during the earlier parts of a trial (Table I) for both interfaces and tasks.³ However, under head array control the temporal distribution is less skewed, indicating that the need for disambiguation

³A uniform temporal distribution corresponds to a trial in which the disambiguation requests are uniformly spread out during the course of task execution. The skewness of a uniform distribution is zero.

TABLE I
SKEWNESS OF THE TEMPORAL DISTRIBUTION OF
DISAMBIGUATION REQUESTS

	Single-step	Multi-step
Joystick	0.63	0.57
Head Array	0.35	0.22

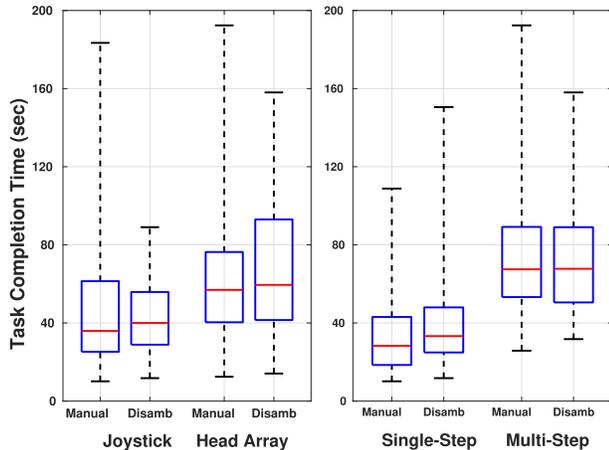


Fig. 7. Task completion times. *Disambiguation* and *Manual* paradigms. Left: Grouped by control interfaces. Right: Grouped by tasks.

request persists throughout the trial, likely due to the extremely low-bandwidth of the interface.

C. Performance

No statistical difference is observed in task completion times between the *Manual* and *Disambiguation* paradigms (Figure 7). However, the variance in the task completion times in general is lower under the *Disambiguation* conditions (with the exception of single-step), indicating more consistent task performance when disambiguation requests are utilized. The task success is 92.50% (148 out of 160 trials) and 97.92% (94 out of 96 trials) for the disambiguation and manual paradigms respectively. 13 out of the 14 unsuccessful trials occur during the more complex multi-step task. Figure 8 provides illustrative examples of the time evolution of goal probabilities. Figure 8 (left) demonstrates how operation in the disambiguating mode can very quickly elevate one goal probability above the threshold for providing autonomy assistance. Figure 8 (right) demonstrates how, at times, subjects do not leverage the capabilities of the disambiguating mode and immediately perform a manual mode switch, without ever issuing any control commands in the disambiguating mode.

D. User Survey

Table II summarizes the results of the user survey. Users agree that task execution is easier during disambiguation trials (Q1, 4.88 ± 0.95) and that operation under disambiguating modes is enjoyable (Q3, 5.00 ± 1.15). User responses strongly validate the effectiveness of the blending-based shared control scheme (Q2, 6.19 ± 0.75). Unsurprisingly, all users feel that it is harder to control the robot using the head array (Q4) and

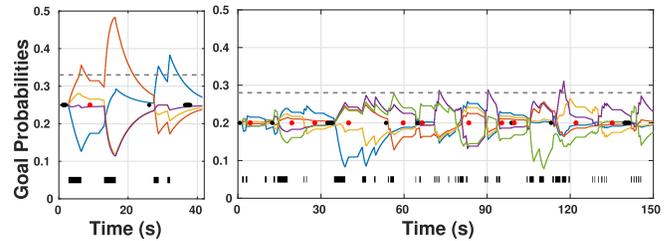


Fig. 8. Time evolution of goal probabilities. Plot annotations include the minimum threshold for robot assistance (gray dashed line), disambiguation requests (red dots), manual mode switches (red dots), and indications of non-zero human control commands (black segments below the plotted data).

TABLE II
SUBJECTIVE SURVEY RESULTS

	Across Tasks	Single-step	Multi-step
Q1	4.88 ± 0.95	4.88 ± 0.99	4.88 ± 0.99
Q2	6.19 ± 0.75	6.25 ± 0.89	6.13 ± 0.64
Q3	5.00 ± 1.15	5.25 ± 1.28	4.75 ± 1.03
Q4	Head Array	Head Array	Head Array
Q5	Head Array	Head Array	Head Array
Q6	Disambiguation	Disambiguation	Disambiguation
Q7	Disamb/Manual	Disamb/Manual	Disamb/Manual

rate the utility value of the disambiguation paradigm to be higher for robot control with the head array (Q5). Although the subjects overwhelmingly prefer the *Disambiguation* to the *Manual* paradigm (Q6) only four out of the eight subjects find the *Disambiguation* paradigm to be user-friendly (Q7). One possible explanation is a lack of transparency regarding why the autonomy chose the disambiguating mode.

VIII. DISCUSSION

The disambiguation algorithm presented in our work can be utilized in any human-robot system in which there is a need to disambiguate between the different states of a discrete hidden variable (for example, a set of discrete goals in robotic manipulation or a set of landmarks in navigation tasks). Our algorithm assumes the existence of a discrete set of parameters (for example, control modes for robotic manipulation or natural language based queries for navigation) that can help the intent inference mechanism to precisely converge to the correct solution. Although the disambiguation algorithm is task-agnostic—because it relies exclusively on the shape features of the probability distribution over the hidden variable—the disambiguation is only as good as the efficacy of the inference algorithm that is used. In our experience, the choice of cost functions and domain-specific heuristics used for inference need to be appropriate for the task at hand. During our implementation development, the efficacy of the disambiguation algorithm degraded when we used only a subset of the four features to inform the disambiguation metric. This only reinforces the need for a combination of different shape features for successful disambiguation.

Another observation from our subject study is how often participants submitted a disambiguation request and then chose not to operate in the selected mode—effectively not letting the robot help them. This under-utilization phenomenon is illustrated in Figure 8 (right). One possible explanation is the

subject's lack of understanding of how their control commands can help the robot to understand their intent. It is likely that a good grasp of the assistance mechanism is critical for providing intent-expressive control commands to the autonomy—underlining the need for extensive and thorough training, and for greater transparency in the human-robot interaction, so that the human has a clear understanding of how and why the autonomy chooses a specific assistance strategy.

The training can be made more effective in a few different ways. For example, the subjects could be explicitly informed of the task relevant features (directedness, proximity *et cetera*) that the autonomy relies on for determining the amount of assistance to offer. Knowledge of these features might motivate the users to leverage the disambiguating mode more.

The inherent time delays associated with the computation of the disambiguation mode (approximately 2-2.5s) might have been a cause for user frustration. Half of the subjects did report that the disambiguation system was not user-friendly. To improve upon this delay, a large set of disambiguating modes could be precomputed for different parts of the workspace, goal configurations and priors ahead of time, which then could be available in a lookup table during task execution. Automated mode switching schemes that eliminate the need for manual button presses altogether might also be a viable option for significantly reducing task effort.

In our future work, as informed by our pilot study, we plan to extend the framework into an automated mode switch assistance system. A more extensive user study with motor-impaired subjects will also be conducted.

IX. CONCLUSION

In this paper, we have presented the idea of *intent disambiguation assistance* via control mode selection for a shared-control assistive robotic arm. The aim of our control mode selection algorithm was to elicit more *intent-expressive* control commands from the user by placing control in those control modes that *maximally disambiguate* between the various goals in a scene. A pilot user study was conducted with eight subjects to evaluate the efficacy of the disambiguation system. Our results indicated a decrease in task effort in terms of the number of button presses when the disambiguation system was active. As our last contribution, we also presented a novel intent inference mechanism inspired by *dynamic field theory* that works in conjunction with the disambiguation system.

ACKNOWLEDGMENT

Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the aforementioned institutions.

REFERENCES

- [1] M. P. LaPlante *et al.*, "Assistive technology devices and home accessibility features: Prevalence, payment, need, and trends," in *Advance Data From Vital and Health Statistics*. Hyattsville, MD, USA: U.S. Department of Health and Human Services, Public Health Service, Centers for Disease Control, National Center for Health Statistics, 1992.
- [2] L. V. Herlant, R. M. Holladay, and S. S. Srinivasa, "Assistive teleoperation of robot arms via automatic time-optimal mode switching," in *Proc. 11th ACM/IEEE Int. Conf. Hum.-Robot Interact. (HRI)*, Mar. 2016, pp. 35–42.
- [3] J. Philips *et al.*, "Adaptive shared control of a brain-actuated simulated wheelchair," in *Proc. IEEE 10th Int. Conf. Rehabil. Robot.*, Jun. 2007, pp. 408–414.
- [4] E. Demeester, A. Hüntemann, D. Vanhooydonck, G. Vanacker, H. Van Brussel, and M. Nuttin, "User-adapted plan recognition and user-adapted shared control: A Bayesian approach to semi-autonomous wheelchair driving," *Auto. Robots*, vol. 24, no. 2, pp. 193–211, Feb. 2008.
- [5] D. Gopinath, S. Jain, and B. D. Argall, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robot. Autom. Lett.*, vol. 2, no. 1, pp. 247–254, Jan. 2017.
- [6] K. Muelling *et al.*, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Auto. Robots*, vol. 41, no. 6, pp. 1401–1422, Aug. 2017.
- [7] P. M. Pilariski, M. R. Dawson, T. Degris, J. P. Carey, and R. S. Sutton, "Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots," in *Proc. 4th IEEE RAS EMBS Int. Conf. Biomed. Robot. Biomechatronics (BioRob)*, Jun. 2012, pp. 296–302.
- [8] C. Liu *et al.*, "Goal inference improves objective and perceived performance in human-robot collaboration," in *Proc. Int. Conf. Auto. Agents Multiagent Syst.*, 2016, pp. 940–948.
- [9] Y. S. Choi, C. D. Anderson, J. D. Glass, and C. C. Kemp, "Laser pointers and a touch screen: Intuitive interfaces for autonomous mobile manipulation for the motor impaired," in *Proc. 10th Int. ACM SIGACCESS Conf. Comput. Accessibility (Assets)*, 2008, pp. 225–232.
- [10] R. Cantrell, M. Scheutz, P. Schermerhorn, and X. Wu, "Robust spoken instruction understanding for HRI," in *Proc. 5th ACM/IEEE Int. Conf. Hum.-Robot Interact. (HRI)*, Mar. 2010, pp. 275–282.
- [11] S. Jain and B. Argall, "Recursive Bayesian human intent recognition in shared-control robotics," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2018, pp. 3905–3912.
- [12] A. D. Dragan and S. S. Srinivasa, *Formalizing Assistive Teleoperation*. Cambridge, MA, USA: MIT, 2012.
- [13] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization for teleoperation and teaming," 2017, *arXiv:1706.00155*. [Online]. Available: <http://arxiv.org/abs/1706.00155>
- [14] H. Admoni and S. Srinivasa, "Predicting user intent through eye gaze for shared autonomy," in *Proc. AAAI Fall Symp. Ser., Shared Autonomy Res. Pract. (AAAI Fall Symp.)*, 2016, pp. 298–303.
- [15] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI*, vol. 8, Chicago, IL, USA, 2008, pp. 1433–1438.
- [16] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *Int. J. Robot. Res.*, vol. 32, no. 7, pp. 790–805, Jun. 2013.
- [17] J. J. Abbott, P. Marayong, and A. M. Okamura, "Haptic virtual fixtures for robot-assisted manipulation," in *Robotics Research*. Berlin, Germany: Springer, 2007, pp. 49–64.
- [18] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *Proc. 8th ACM/IEEE Int. Conf. Hum.-Robot Interact. (HRI)*, Mar. 2013, pp. 301–308.
- [19] R. M. Holladay, A. D. Dragan, and S. S. Srinivasa, "Legible robot pointing," in *Proc. 23rd IEEE Int. Symp. Robot Hum. Interact. Commun.*, Aug. 2014, pp. 217–223.
- [20] D. Gopinath and B. Argall, "Mode switch assistance to maximize human intent disambiguation," in *Proc. Robot., Sci. Syst.*, 2017.
- [21] G. Schöner and J. Spencer, *Dynamic Thinking: A Primer on Dynamic Field Theory*. Oxford, U.K.: Oxford Univ. Press, 2015.
- [22] S.-I. Amari, "Dynamics of pattern formation in lateral-inhibition type neural fields," *Biol. Cybern.*, vol. 27, no. 2, pp. 77–87, 1977.
- [23] G. Schöner, "The dynamical systems approach to cognition," in *Cambridge Handbook of Computational Cognitive Modeling*. Cambridge, U.K.: Cambridge Univ. Press, 2008, pp. 101–126.
- [24] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robot. Auto. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.
- [25] D. Hsu, R. Kindel, J.-C. Latombe, and S. Rock, "Randomized kinodynamic motion planning with moving obstacles," *Int. J. Robot. Res.*, vol. 21, no. 3, pp. 233–255, Mar. 2002.
- [26] E. Rimon and D. E. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Trans. Robot. Autom.*, vol. 8, no. 5, pp. 501–518, Oct. 1992.
- [27] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. Robot. Res.*, vol. 5, no. 1, pp. 90–98, 1986.