Learning Control Sharing Strategies for Assistive Cyber-Physical Systems

PI: Brenna Argall, Northwestern University

[Mutli-PI project with Siddhartha Srinivasa, University of Washington]

NSF Award No: 1544741

Why Do We Need Shared Autonomy in Assistive Cyber-Physical Systems?

Figure 1: ***Top:*** **7-DoF robotic arm with 2D joystick and 1D head array control interfaces** ***Bottom:*** **Example partitioning of 7-D control space into operational control modes.**

Aim: To develop principled approaches for shared control of complex assistive cyber-physical systems, such as a 6-DoF robotic arms via simple low-dimensional control interfaces that are accessible to persons with severe motor impairments such as 2D joysticks and 1D Sip-N-Puff interfaces (Figure 1, Top).

Why: The dimensionality mismatch between high-dimensional robots and low-dimensional control interfaces requires the control space to be partitioned into *control modes* (Figure 1, Bottom). For full control of the robot the user switches between these partitions and this is known as *mode switching.* Mode switching adds to the cognitive workload and degrades task performance. *Shared autonomy* helps to alleviate some of the task burden by letting the robot take partial responsibility of task execution.

In this project, we innovate ways in which *automated mode-switching assistance* can help to alleviate the control burden of operating assistive machines.

Relevance to Cyber-Physical Systems

Synergy: The intersection of robotic manipulation, human rehabilitation, control theory, machine learning, human-robot interaction and clinical studies.

Science of CPS: Development of mathematical models to quantify the interaction dynamics between the user and system. Development of new interfaces and interaction modalities with strong theoretical foundations.

Engineering of CPS: Deployment of algorithms on real hardware and evaluation of proposed assistive paradigms with able-bodied and spinal cord injured users.

Dynamic Neural Fields for Inference

$$\frac{\partial \boldsymbol{p}(t)}{\partial t} = \frac{1}{\tau}\Big[\underbrace{-\boldsymbol{P}^T_{n_g\times n_g}\cdot \boldsymbol{p}(t)}_{\text{goal transition dynamics}} + \underbrace{\frac{1}{n_g}\cdot \mathbb{1}_{n_g}}_{\text{rest state}}\Big] + \underbrace{\boldsymbol{\lambda}_{n_g\times n_g}\cdot\sigma(\boldsymbol{\xi}(\boldsymbol{u}_h;\boldsymbol{\Theta}))}_{\text{excitatory + inhibitory}}$$

Idea: Use a *dynamic neural fields* based approach to specify the time-evolution of the probability distribution over goals as a dynamical system with constraints.

Why: The power of a disambiguation algorithm is closely linked to the accuracy of the inference scheme. By casting belief propagation as a dynamical system with constraints, information from past actions can be easily incorporated.

Results: Our field-theoretic inference approach outperformed memory-based prediction significantly (87.46% vs. 59.15%, respectively) and was comparable to recursive belief updating (87.43%) [1]. An illustrative example in Figure 4 shows the goal probability evolutions in the absence of information (no control commands).

[1] Deepak Gopinath and Brenna D. Argall. *A Framework for Goal Disambiguation for Assistive Robotics.* Under review.

[2] Deepak Gopinath and Brenna D. Argall, *Information-Theoretic Characterization of Control Modes for Intent Disambiguation in Assistive Robots.* Under review.

Information Theoretic Approaches to Intent Disambiguation

Aim: To elicit more *intent-expressive* control commands from the human co-operator by identifying and autonomously switching into control modes within which operation maximally disambiguates their intent.

Why: The general idea is one of "help me help you"---whereby placing the system into a control mode which better disambiguates human intent during teleoperation, the human more quickly reveals his/her intent to the autonomy. The autonomy as a result is able to step in to provide control assistance sooner and with greater accuracy. Our approach uses the information-theoretic principles of entropy gain and KL-divergence to characterize information gain in regards to the probability distribution over goals (which represents user intent) that will result from user-initiated robot motion along each control dimension.

Results: We performed simulation-based experiments on point robots operating in task spaces of different dimensions (R2, SE(2), R3, SE(3)) and a simulated robotic arm operating in SE(3) using a 1D discrete interface [2]. Four different disambiguation metrics were tested (entropy gain, KL-divergence, a heuristic, and greedy potential maximization (baseline) within a mode-switch assistance paradigm, as well as three different intent inference schemes (heuristic, Bayesian, and field-theoretic) and .Disambiguation enabled the autonomy to provide assistance *earlier* in the task execution, for all tested disambiguation metrics (Figures 2 and 3). Furthermore, at least one disambiguation metric always resulted in an intent inference *accuracy increase* compared to the baseline.

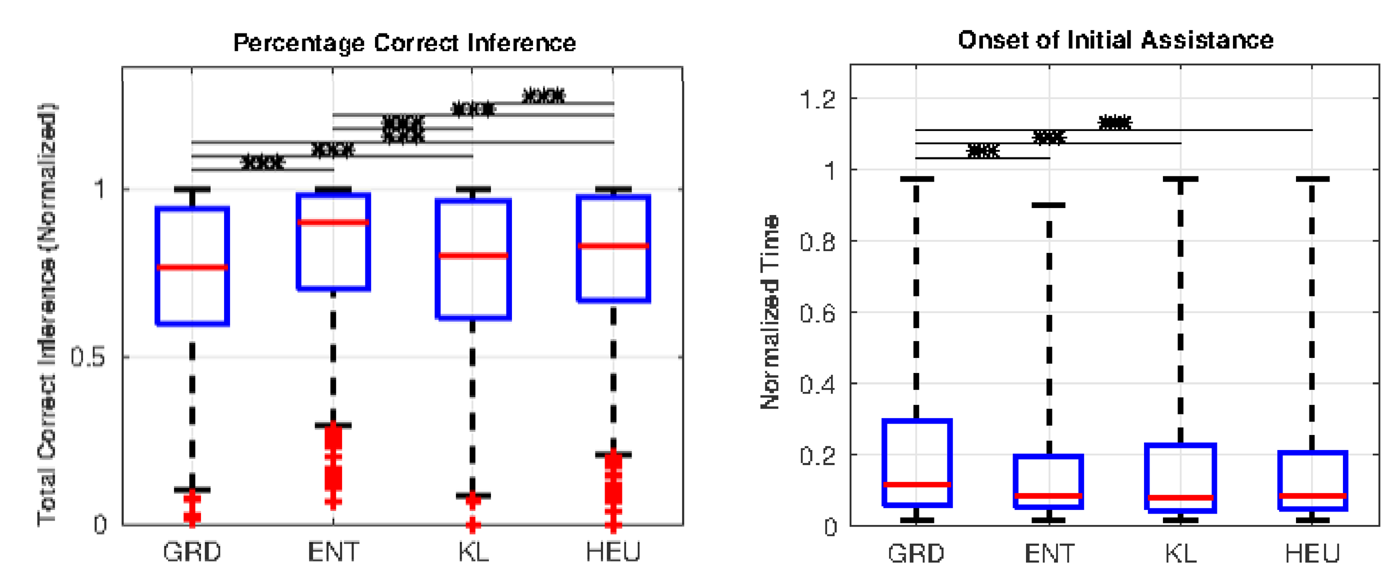

Figure 2: Disambiguation assistance results with a simulated point robot in R3. Comparison of three disambiguation metrics (two information-theoretic, ENT and KL, one heuristic-based, HEU) and a greedy baseline (GRD). ***: p<0.001 (Wilcoxon Rank-Sum test)

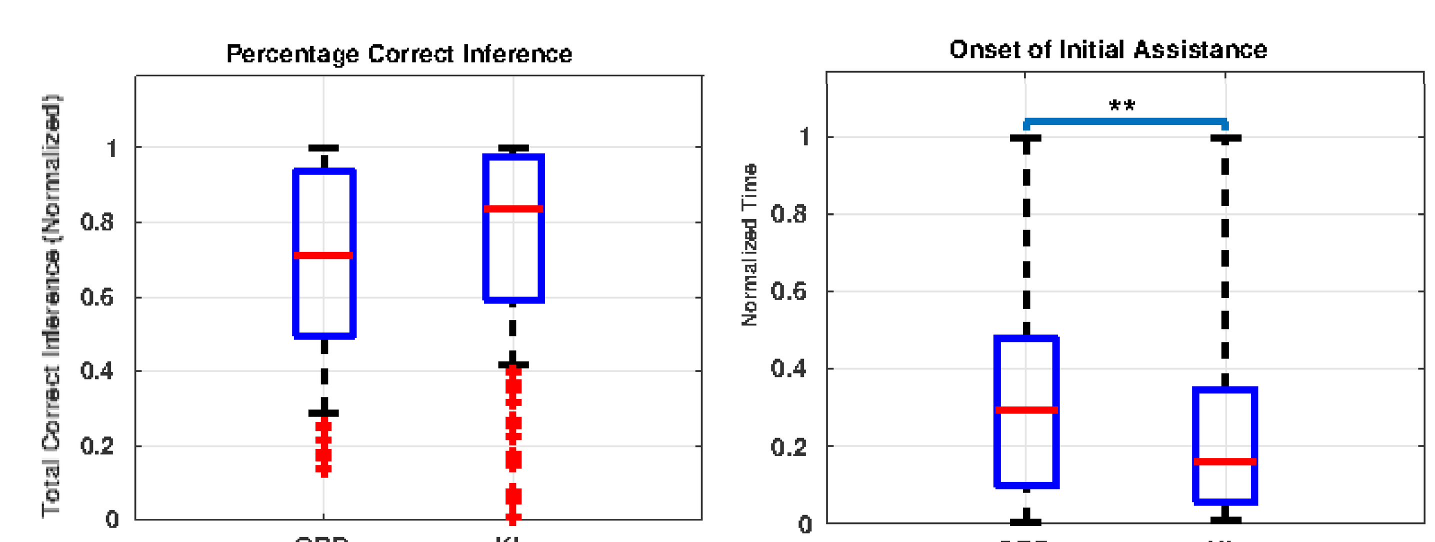

Figure 3: Disambiguation assistance results with a simulated 6-DoF robotic arm. Comparison of our disambiguation approach (KL) and the greedy baseline (GRD). **: p<0.01 (Wilcoxon Rank-Sum test)

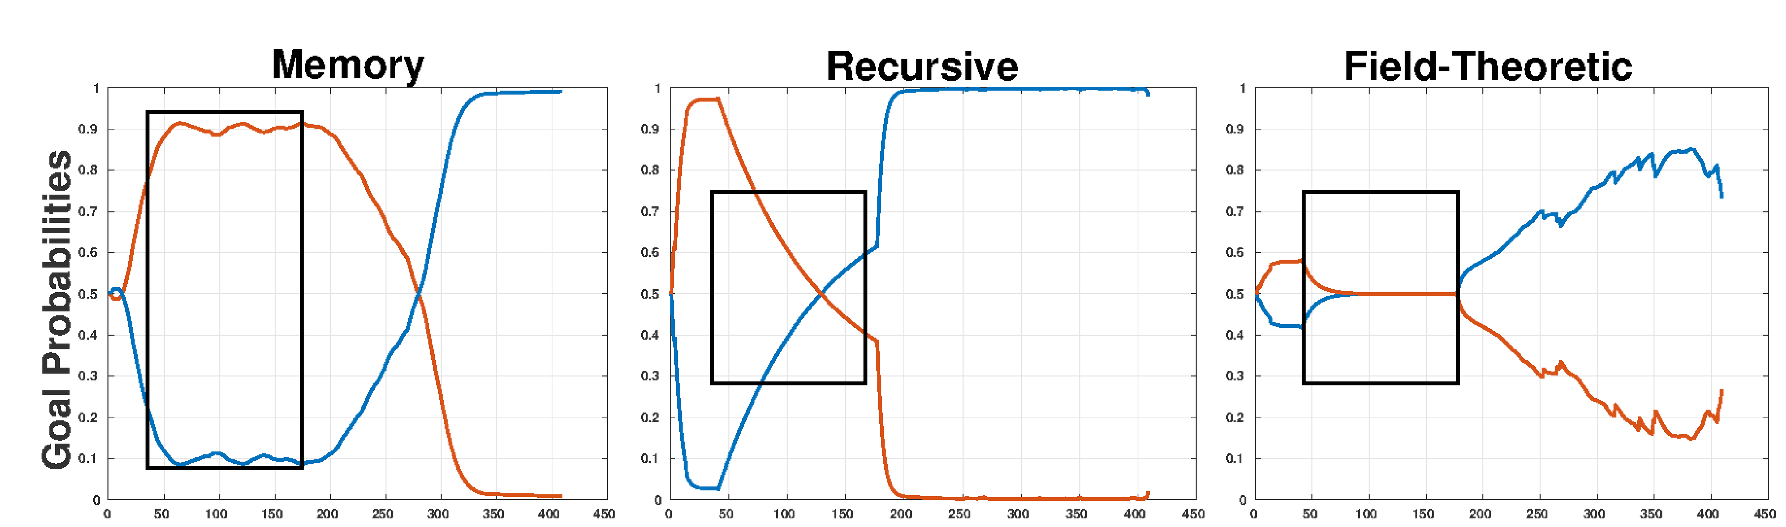

Figure 4: Comparison of inference approaches during human teleoperation to one of two goals (orange or blue). Under zero velocity commands (black box outline), only our field-theoretic approach correctly converges to the maximum entropy uniform distribution.