Inverted Trust Improves Shared Control of Complex Dynamic Systems

Alexander Broad*, Matthew Derry*, Jarvis Schultz*, Todd Murphey* and Brenna Argall*†
* Northwestern University, Evanston, Illinois 60628
† Rehabilitation Institute of Chicago, Chicago, Illinois, USA
Email: {alex.broad,matthewderry2011}@u.northwestern.edu, {jschultz,t-murphey,brenna.argall}@northwestern.edu

I. Introduction

The increasing pervasiveness, capabilities and complexity of autonomous robots in human environments has highlighted the need for more sophisticated control-sharing techniques that allow humans to interact with, control and shape the behaviors of these systems, while also maintaining a high level of safety [2]. Control sharing can create a system that leverages the strengths of each source of control while reducing the effects of the weaknesses [1]. Our premise is that an understanding of each by the other is essential for successful shared control systems. We quantify human understanding of the dynamic system through a metric of trust.

Our take on control sharing is to combine the relative advantages of the human and robotic partners. In this work, we consider that automation and optimal control techniques are good at controlling highly dynamic systems, but require a reference trajectory to try to stabilize to. While these reference trajectories could come from automated path planners, engaging a human partner has the advantage of using the exceptional perceptual capabilities and situational awareness that humans often have to operate in dynamic environments. The key is for the human to provide reference trajectories that the automated controller can track. We develop a human-in-the-loop control framework that reasons explicitly about the amount of control authority that should be allocated to the human based on the trust that the autonomous system has *in the human*. The purpose of this trust metric is to allow the system to learn how able a user is in providing reference trajectories that can be easily tracked by the automated controller. The adaptive trust metric can then be used to develop a safer and more stable shared-control system.

II. Formulation of Trust

In order to define, adapt and make use of a formal notion of trust, our proposed framework consists of two steps. In the *evaluation of user input* step, a trust metric is calculated as a function of the deviation from the reference trajectory. The *control modulation* step then uses the trust metric to allocate control authority.

After each interaction with the system, a trust metric is calculated using tools provided by optimal control theory. We take a control-theoretic viewpoint in which we use the *deviation* of the executed trajectory from the reference trajectory, as this measure indicates how well the receding horizon controller is able to track the user input. For a given trial i we calculate a deviation metric δ^i. The deviation from the reference trajectory can be captured by the Fréchet distance between the executed and desired trajectories,

$$\delta^i(f,g) = \inf_{\alpha,\beta} \max_{t\in[0,1]} d(f(\alpha(t)), g(\beta(t))) \tag{1}$$

where $f : [a,b] \to V$ and $a,b \in \Re$, is the reference trajectory and $g : [a',b'] \to V$ and $a',b' \in \Re$, is the control trajectory and (V,d) is a metric space. α and β are continuous nondecreasing functions that map from $[0,1]$ onto $[a,b]$ and $[a',b']$, respectively.

The trust metric should decrease as deviation δ increases. To calculate the trust metric τ^i at trial i, we represent the distribution of deviations as a Gaussian distribution, $\delta \sim N(\mu,\sigma^2)$, where μ and σ^2 are the mean and variance of an individual's deviation history. We then update the previous trust metric τ^{i-1} by computing the probability $\mathscr{P}$ of δ^i

$$\tau^i = \tau^{i-1} - \gamma \cdot \mathscr{P}\{\delta = \delta^i\} \tag{2}$$

where γ is a learning rate that determines how quickly the trust decays with performance, and $0.1 \leq \tau \leq 1$.

This Gaussian distribution represents the system's knowledge of the user, and is iteratively updated after each trial. Intuitively, if there is a large deviation from the desired trajectory, then τ^i should be low-because the user is providing reference trajectories that are difficult for the controller to track, which reduces the overall robustness of the system.

Modulation of the user's input is realized through a combination of a low-pass filter and scaling the input speed. By removing the high-frequency content from the input signal, the receding horizon controller is better able to track the reference trajectory. Similarly, by scaling the input speed, users are better able to control for momentary mistakes that can lead to challenging reference trajectories. It is important to note that these transformations can adversely affect more typical task performance measures such as time to completion, but this is a trade-off for system safety and stability.

III. Experiment Design and Result

Our trust-based shared-autonomy framework is demonstrated on a simulated planar crane system, in which an overhead robot with a winch has a mass suspended by a string. We choose this system because it provides dynamics that are difficult for a human to control, while allowing for the definition of dynamic tasks that are representative of tasks for which control sharing would be beneficial. Here we present selected results from this study.

A. Experimental Setup

Inspired by a task that is currently part of real crane operator certification tests, we implement a maze navigation task within our simulated planar crane environment. We test three different task configurations of increasing difficulty. First, a *low* difficulty configuration where the total path length is short, requires few turns ($\sim$3) and the maze hallways are wide. Then, a *medium* difficulty configuration where the total path length is longer, requires more turns ($\sim$5) and the maze hallways are an average of 60% more narrow then the low difficulty configuration. Finally, a *high* difficulty configuration where the total path length is long, includes many turns ($\sim$10) and the maze hallways are an average of 50% more narrow than the low difficulty configuration. A total of 21 (8 low difficulty, 8 medium difficulty and 5 high difficulty) unique mazes are used in this experiment.

Twenty-two users were recruited from the Northwestern University community.[1] Users were randomly grouped into two cohorts. Static: The trust level was held static after the initial training period, and Adaptive: The trust level evolved throughout the experimental procedure. After the trust initialization, each experiment consisted in a total of 20 trials. Five trials of direct control on the low-difficulty maze and five trials of shared-control on the each of the 3 levels of maze difficulty.

B. Results

We analyzed the system's ability to modulate a user's trust metric to produce safer and more stable reference trajectories. We performed a statistical analysis comparing the average controller magnitude between the Adaptive and Static trust cohorts in each maze configuration.

Statistical analysis was done using a two-tailed Student's t-test. In both the static ($p < 0.01$) and adaptive ($p < 0.01$) trust cohorts, we saw a statistically significant decrease in the average controller magnitude required to track the user's reference trajectory in the final maze configuration when compared with the initial maze configuration. This suggests that users in both cohorts are able to learn pertinent aspects of the system dynamics and how to provide safe and stable reference trajectories from the viewpoint of the automated controller.

We also found (Fig. 1) a statistically significant difference between the average controller magnitude required to track references trajectories provided by users in the static trust cohort versus those in the adaptive trust cohort, in the final maze configuration ($p < 0.01$). As we see no statistical evidence that one cohort outperforms the other in the first two maze configurations, we can infer that the adaptive trust formulation allows the system to adapt to the (possibly changing) abilities of the user, and so modulate the user input to provide safer reference trajectories for the controller to stabilize to—regardless of the abilities of the individual user.

[1]Two participants were removed from the study before analysis due to a poor understanding of the task requirements. On average, users completed 11.7 of the 15 test trials. The two removed users completed one trial each.

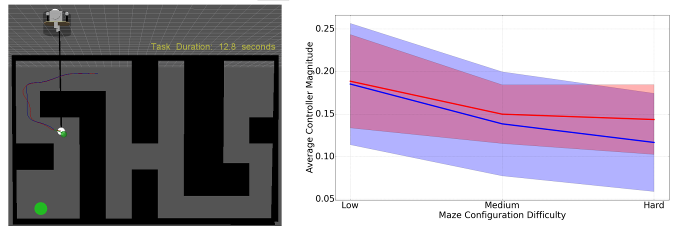

Fig. 1. Left: Simulated crane environment. Right: Average controller magnitude per maze configuration for the static (red) and adaptive (blue) trust cohorts. Mean (line) and standard deviation (variance envelope) are presented.

IV. Conclusion and Future Work

Thus, results show that an adaptive trust metric, based on our control-theoretic formulation, was able to improve the ability of the shared-control system to produce references trajectories that were significantly ($p < 0.01$) easier for the controller to track than those provided by users with a static trust metric. The reduced average controller magnitude reflects the system's ability to learn appropriate methods for modulating the operator's input, resulting in reference trajectories that are easier to track. This work creates a foundation upon which to expand the trust-based shared control framework to include the online, continuous adaptation of trust, more granular user skill level classification, as well as applications to additional tasks and robot platforms.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant CNS 1329891. 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] R. Chipalkatty, H. Daepp, M. Egerstedt, and W. Book. Human-in-the-loop: MPC for shared control of a quadruped rescue robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4556–4561, 2011.

[2] Kakin K. Tsoi, Mark Mulder, and David A. Abbink. Balancing safety and support: Changing lanes with a haptic lane-keeping support system. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1236–1243. IEEE, 2010.