On-line Human Intervention for Robot Behavior Adaptation

Maggie Wigness[1] and John G. Rogers III[1]

***Abstract*— Unmanned ground vehicles operating in unstructured environments as part of a human-robot team should be expected to learn and adapt on-line during operation just as their human counterparts. Further, this on-line adaptation should require minimal human teammate oversight to ensure each teammate can maintain an appropriate level of awareness. We discuss how human intervention can be used during on-line operation to quickly correct and adapt navigation behaviors learned from inverse reinforcement learning. We discuss the relevance of adaptation with respect to long-term environment changes, domain shifts, and novel information discovery.**

I. INTRODUCTION

When human-robot teams operate together in a dynamic environment, it can be expected that behaviors of each teammate will need to adapt based on the current environment and mission status. For example, severe weather or seasonal shifts could drastically change the terrain, causing areas that were previously traversable to present safety issues. Further, novel information, terrain or objects may be encountered throughout operation that require adaptation.

In our previous work [1], we described how to use inverse reinforcement learning to quickly train reward functions that encode traversal behaviors, e.g., driving near the edge. Our approach was based on the principal of maximum entropy [2], and was shown to learn reliable reward functions with only a few human demonstrations. Although this learning technique required minimal human effort, the training process was performed completely off-line through batch processing a set of demonstrated trajectories.

For human-robot teams to collaborate at high operational tempo, it will be necessary for the robot teammates to adapt on-line when their behaviors deteriorate due to changes in the environment or mission requirements. We discuss how our previous learning from demonstration technique can be modified for learning from intervention. In this case, the human teammate has the ability to interrupt autonomous robot operation at any moment to provide corrected traversal behavior. Human interventions allow the robot to gain additional training exemplars without being pulled from operation, and instead, use the trajectories provided during intervention to directly update the behavior model with onboard processing.

II. RELATED WORK

Learning from intervention has been used for a variety of tasks. Saunders et al. [3] discuss the utility of learning from intervention and its ability to reduce catastrophic failures when intelligent systems are learning. The combination of demonstration and intervention has been explored in simulation for an autonomous perching task [4], and for robot manipulation and gesture tasks [5]. We use similar human intervention styles as these approaches, but focus on adapting traversal behaviors on a mobile robot in a real-world environment.

[1]U.S. Army Research Laboratory, Adelphi, MD, USA {maggie.b.wigness,john.g.rogers59}.civ@mail.mil

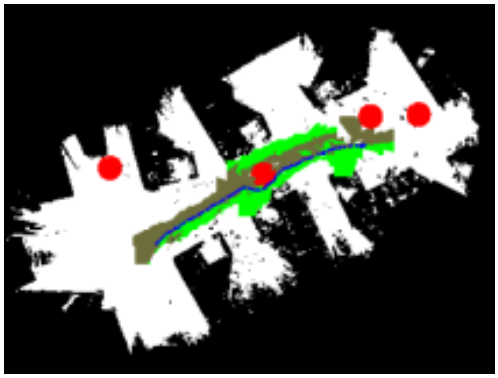

Fig. 1. Illustration of four environment feature maps used for the inverse reinforcement learning of traversal behaviors. Human demonstration of "optimal" traversal behavior (denoted by the blue line) is performed with respect to the features obstacles (black), grass (green), road (gold), and dangerous regions (red).

III. EXPERIMENTS

We use our previous learning from demonstration setup [1] that learns a reward function given 1) a set of demonstrated trajectories that exemplify "optimal" traversal patterns for the desired behavior, and 2) a set of environment feature maps, i.e., binary occupancy grids that encode the presence or absence of specific terrain and objects. Fig. 1 shows a color-coded overlay of four environment feature types used in the behavior learning: obstacles (black), grass (green), road (gold), and dangerous regions (red).

We define an experimental scenario in which the robot has been deployed to execute a task with a traversal behavior reward function that was learned from human demonstration off-line. This reward function encodes the "edge of road" traversal behavior, where the robot should navigate throughout the environment while trying to maintain close proximity to the road's edge. During operation, the robot's behavior needs to be adapted as additional environment information is encountered. Specifically, the human-robot team is given additional information about potentially dangerous regions in the environment. For experimentation, these dangerous areas are encoded as circular regions with a pre-defined radius.

Fig. 2 depicts the traversal behavior planned by the robot in the presence of these dangerous regions given its original