

Robot Learning With Predictions of Operator Intent*

Michael T. Rosenstein, Andrew H. Fagg and Roderic A. Grupen

Department of Computer Science
University of Massachusetts Amherst
140 Governors Drive
Amherst, MA 01003
{mtr,fagg,gruppen}@cs.umass.edu

Abstract

The goal of this paper is to sketch a scenario for robot learning with remote operation by a person. We are particularly interested in situations where a user interface provides not only feedback about robot performance but also predictions of operator intent. We describe several benefits and challenges of this learning paradigm.

Introduction

Teleoperation, i.e., human control of a machine across some barrier (Johnsen & Corliss 1971), remains perhaps the most effective means for solving many robotics tasks beyond the factory floor. Examples include hazardous waste disposal, planetary and undersea exploration, robot-assisted surgery, and emerging applications such as elder care. In contrast to factory automation, the key feature of these tasks—and the key challenge for researchers to deal with—is the variation in circumstances from one instance of the task to the next. This makes general-purpose, automated solutions difficult to devise, even for a restricted class of tasks. One advantage of teleoperation over automatic control is that a human operator remains in the robot’s control loop to provide intelligent decision-making, thereby mitigating the need for an intelligent robot. As a long-term solution, however, current forms of teleoperation suffer from the person’s inability to control a robot as easily as his or her own limbs. This leads to movements and “cognitive load” that are fatiguing for the operator and that fail to utilize the robot’s full potential, especially in terms of speed.

One remedy for these difficulties is the use of *adjustable autonomy*, e.g., (Kortenkamp, Schreckenghost, & Bonasso 2000; Scerri, Pynadath, & Tambe 2002), whereby a computer interface facilitates switching among a variety of control modes ranging from full supervision to full autonomy. As robots improve their cognitive capabilities—for perception, reasoning, and action—such interfaces will give way to modes of interaction where the robot plays the role of partner rather than tool. In the meantime, one goal of human-

robot interface design is to shift some responsibility, i.e., cognitive load, from human to machine.

As a recent example, Fong et al. (2003) proposed a dialogue-based interface where human and robot exchange messages using a simple language. Operators used a hand-held computer to perform precision driving of a mobile robot and also to give up control when the autonomous system sent queries such as “Can I drive through?” Autonomous behavior then gave the operator freedom to supervise a second mobile robot (Fong, Thorpe, & Baur 2003).

Examples such as this involve several steps executed by the interface software: (1) Gather information about the current state of affairs. (2) Use that information to predict the operator’s intentions. (3) If those intentions match the capabilities of the autonomous system, then indicate to the operator that this is the case. (4) Allow the user to initiate a switch to automatic control. (5) Indicate termination of the current sub-task and return control to the operator, or else allow the operator to interrupt the autonomous system.

These steps are repeated for the duration of the human-robot interaction, although considerable variation is possible. For example, the interface may not actually predict the operator’s intentions but rather present all admissible commands sorted by some pre-specified criterion. Similarly, in some pre-determined situations, the interface may turn over control to the autonomous system without waiting for operator confirmation.

In addition to the potential for solving real-world tasks, teleoperation with adjustable autonomy also provides a wealth of information from which a robot could learn. Presumably the human operator has some expertise in the task at hand, and so the robot has the opportunity to bootstrap its knowledge and skills by associating its own sensory information with commands from the operator. For instance, if expert operators always approach an object from a particular direction (relative to the object) then this may be useful evidence to bias a machine vision system watching for salient image features. Similarly, a sequence of steps executed by the operator could be used to learn precedence conditions for an assembly task. A robot could also develop fault-tolerant behavior by learning from situations where a person interrupts to perform a low-level maintenance operation, e.g., obstacle avoidance. In any case, the ability to learn from experience is likely to be a critical component of

*This work was supported by NASA under award No. NAG-9-1379. 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 National Aeronautics and Space Administration.




Figure 1: A tool-use experiment with (a) Dexter and (b) the corresponding predictor display. Candidate landmarks are highlighted with red circles, the size of which indicate the ranked predictions of user intent.

future robotic systems.

Predictor Displays

Predictor displays are one of the key interface elements for teleoperation—and especially for “human-in-the-loop” learning. As defined by Sheridan (1992) predictor displays are “cursors or other indications driven by a computer [that] are extrapolated forward in time....” Essentially the computer displays the output from a forward model of the relevant task dynamics, while possibly accounting for future operator inputs as well. Predictor displays are typically used to overcome communication delays as is common with space exploration (Kim & Bejczy 1993) although such displays are also useful for predicting operator intentions.

Figure 1 shows the predictor display for a tool-use experiment with Dexter, a humanoid robot developed at UMass Amherst. For this experiment, predictions of operator intent are based on movement of the robot in the vicinity of key landmarks, i.e., bolts in the figure. In particular, the predictions about which bolt to tighten are simply a function of the position of the end-effector, a cordless drill, relative to each bolt. The operator controls the drill using hand movements with a glove-like input device.

With the proper feedback from the predictor display, the operator has an effective means for adjusting his or her commands *to elicit the desired prediction from the interface*. In other words, the operator has two strategies for tightening a bolt: (1) carefully move the robot toward the target bolt and then engage an autonomous docking controller; or (2) exploit the interface by first moving to a region of space where the target bolt is easily discriminated from the others, and then engage the docking controller. Our anecdotal evidence shows that the latter strategy is extremely effective, and in many cases the operator can use gross movements for precise target selection. (Future work includes a user study to quantify the observed benefits.)

Future Challenges

For control purposes, the advantage of predictor displays is straightforward; they help the operator anticipate future task requirements. But for learning purposes, predictor displays

can be used in two fundamentally different ways: either for robot learning from demonstration or for *simultaneous* learning by human and robot. In both cases, the operator can use the predictor display to decide if the autonomous system needs to be interrupted, such as when predicted movement of the end-effector is toward the wrong target. Under such conditions, one possibility for learning is that the operator then takes control and demonstrates the correct movement, and supervised learning methods generalize the demonstrated “skill.”

The other possibility suggested above is that the operator can try to elicit the desired prediction from the interface. Then the potential for machine learning is to make the operator’s job easier by modifying the mapping from end-effector movement to predicted intent. Essentially, the interface should learn virtual landmarks that are easier to reach than the corresponding real landmarks. The confounding factor is that predictions of intent involve biases that a human operator can learn to exploit, possibly unintentionally. Thus, one challenge for human-in-the-loop machine learning is to deal with the nonstationarities introduced by the operator’s changing behavior.

Another challenge for learning with predictions of operator intent is to balance the benefits and drawbacks associated with the two learning problems (i.e., learning to control the robot versus learning to improve the interface). A third challenge not discussed in this paper is how to design an interface that displays enough information for the operator to trust the autonomous system but not so much information as to overwhelm the operator. Despite these challenges, teleoperation holds great potential for improving a robot’s autonomous capabilities at the same time as it enables real-world tasks to be solved.

References

- Fong, T.; Thorpe, C.; and Baur, C. 2003. Multi-robot remote driving with collaborative control. *IEEE Transactions on Industrial Electronics* 50(4):699–704.
- Johnsen, E. G., and Corliss, W. R. 1971. *Human Factors Applications In Teleoperator Design and Operations*. New York: John Wiley & Sons, Inc.
- Kim, W. S., and Bejczy, A. K. 1993. Demonstration of a high-fidelity predictive/preview display technique for telerobotic servicing in space. *IEEE Transactions on Robotics and Automation* 9(5):698–702.
- Kortenkamp, D.; Schreckenghost, D.; and Bonasso, R. P. 2000. Adjustable control autonomy for manned space flight. In *IEEE Aerospace Conference Proceedings*, 18–25. Piscataway, NJ: IEEE.
- Scerri, P.; Pynadath, D. V.; and Tambe, M. 2002. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research* 17:171–228.
- Sheridan, T. B. 1992. *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: The MIT Press.