SCHooL: Scalable Collaborative Human–Robot Learning

Ken Golderg, Pieter Abbeel, Anca Dragan, Stuart Russell

contact: <goldberg@berkeley.edu> University of California, Berkeley schoolproject.berkeley.edu

1. Overview

- Learning from Demonstrations (LfD) paradigms lack a theoretical framework of human–robot cooperative learning and and hierarchical planning.
- The SCHooL project aims to fill this gap by investigating scalable robot manipulation, where multiple robots collaboratively learn from multiple humans with a unified game-theoretic inverse reinforcement learning framework.
- Integrative Application: "Surface Decluttering": robots that keep specified surfaces clear by identifying, grasping, and appropriately relocating objects with applications in homes, schools, warehouses, offices, manufacturing and machine shops, retail stores using an emerging class of mobile manipulators such as the Fetch robot.

2. Research Objectives

1. Formalize framework for scalable collaborative inverse reinforcement learning (SCIRL) using theory of multiagent games and collaborative learning in multiple distributed domains.
2. Develop deep learning representations of visuospatial features and reward functions to extract and share deep learning representations for scalable human-robot learning.
3. Develop hierarchical task and reward structure to increase planning horizon and decrease sample complexity by partitioning complex tasks into sub tasks.
4. Develop new models to represent and share awareness of robot capabilities and robot models of human intent to support distributed learning.

These objectives support the NRI 2.0 themes: Collaboration, Interaction, and Scalability and for Broader Outreach we are partnering with Lawrence Hall of Science and Penguin Books.

3. Problem Formulation

- Cooperative inverse reinforcement learning (CIRL) is posed as a 2-player cooperative Markov game with asymmetric information: only the human observes the reward parameters, and not the robot.

$$\mathrm{G} = \langle \mathcal{S}, \{\mathcal{A}^{\mathbf{H}}, \mathcal{A}^{\mathbf{R}}\}, P(\cdot|\cdot,\cdot,\cdot), P_0(\cdot), R_\theta(\cdot,\cdot,\cdot), \gamma \rangle$$

- The goal of the robot is to learn the reward parameters by querying the human and the observations of the other agent's actions.
- The robot learns the parameters incrementally by learning from human actions and the human incrementally learns to convey the task to the robot.

4. Primary Results

- A reformulation of AI replacing the standard model (optimizing a fixed, known objective) with optimizing human objectives that are not fully observed.
- New Formal models exploring irrationality in reward inference; learning to control a fleet of robots by humans; learning efficient representation for intrinsic motivation.
- Combining depth sensing and sim-to-real transfer for extracting hierarchical task and reward structure.

Broader Impacts:

- Project incorporated into 4th Edition of Russel and Norvig: **AI: A Modern Approach**. 2020 texbook.
- **How To Train Your Robot.** Elementary school book based on the project to inspire young readers (esp. under-represented minorities) to explore AI and robot learning. Freely distributing 1800 copies to schools and student clubs, featured in IEEE Robot Gift Guide.
-

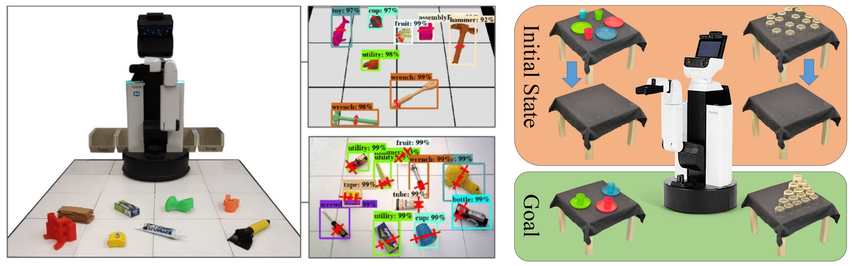

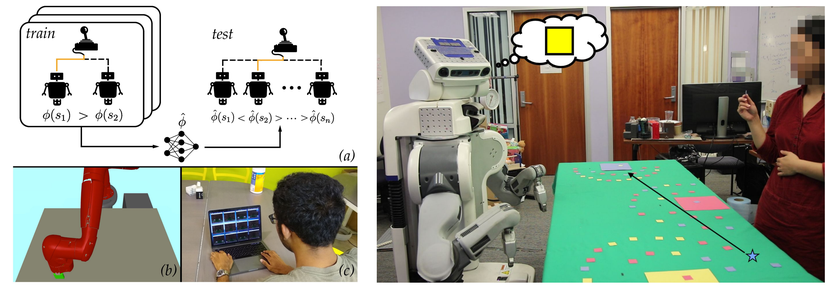

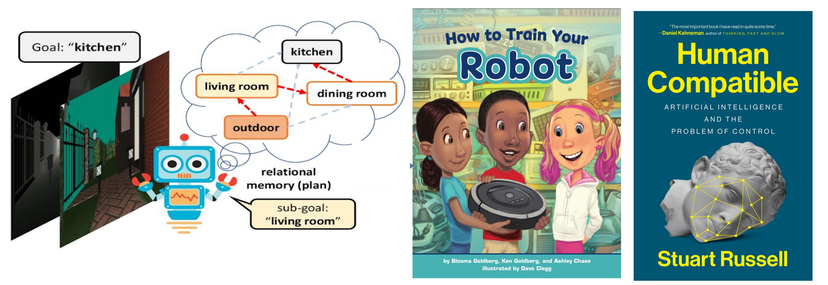

5. Selected Publications

See website above for comprehensive list of over 30 papers.

- G. Swamy, S. Reddy, S. Levine, A. Dragan, "Scaled Autonomy: Enabling Human Operators to Control Robot Fleets", ICRA 2020
- A. Tanwani, P. Sermanet, A. Yan, R. Anand, M. Phielipp, K. Goldberg, "Motion2Vec representation learning from surgical videos", ICRA 2020
- Y. Du, S. Tiomkin, E. Kiciman, A. Dragan, P. Abbeel, "Goal Agnostic Assistance through Human Empowerment", ICML, 2020 (submitted)
- Y. Wu, W. Wu, A. Tamar, S. Russel, G. Gkioxari, Y. Tian, "Bayesian Relational Memory for Semantic Visual Navigation", ICCV, 2019
- I. Huang, S. Huang, R. Pandya, A. Dragan, "Nonverbal Feedback for Human Teachers", CoRL, 2019
- S. Russel, Human Compatible: AI and the Problem of Control. Penguin Books Ltd., 2019
- B. Goldberg, K. Goldberg, A. Chase. How to Train Your Robot. Lawrence Hall of Science, 2019.
- R. Fox, R. Berenstein, I. Stoica, K. Goldberg, "Multi-Task Hierarchical Imitation Learning for Home Automation", CASE 2019