NRI: INT: COLLAB: Collaborative Task Planning and Learning through Language Communication in a Human-Robot Team

IIS 1949634 (formerly 1830244): Joyce Chai, Department of EECS, University of Michigan
IIS 1830282: Julie Shah, Department of Aeronautics and Astronautics, Computer Science and AI Lab, Massachusetts Institute of Technology

Motivation and Overall Objectives

Challenges

- Most robotic planning systems operate based on a closed set of representations. However, in reality robots are not likely to have a complete set of domain models and will often encounter new situations or new tasks they don't have any knowledge or experience about.
- Even given sufficient state and action representations, automated planning is still computationally challenging. Designing planners that can generate high quality plans and perform efficiently across various domains remains an open problem.

Overall Objective

Develop a novel framework that tightly links language and dialogue processing with the robot's underlying planning system to support collaborative task planning and learning in a human-robot team.

Style Learning for Representation and Communication

Can an agent learn how other agents behave, enabling efficient teamwork?

In teams, agents adopt social conventions like driving on the left or right side of the road, or speaking English or French, but current techniques for learning conventions require paired data that are hard to acquire. Using our technique of adversarial pruning, our agents learn to behave according to others' social conventions.

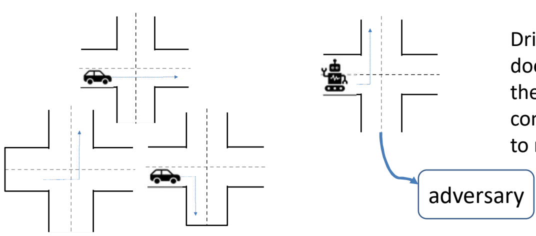

Driving on the left doesn't look like the social convention, so train to not do that.

In assessing mixed-team performance – how well teams with some pre-trained agents and some of our agents – we outperform other methods despite using less paired data.

- In a coordinated lever-pulling game, agents communicate with vectors (2D projection colored here by agent ID). A base model (left) learns a particular sorting order; our model (right) learns that same ordering with just 8 examples.

- Our MNIST autoencoder (right) learns the same latent space, colored here by digit, as a base model (left) after 8 examples.

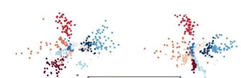

- Speaker-actor agent pairs coordinate to move an actor in blue towards a target in red based on the speaker's commands. Our speakers pair well with pre-trained actors, and vice versa.

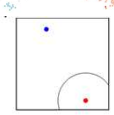

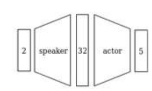

Extensions of this work could include learning human social conventions using similar techniques.

Action-Effect Learning

Action to Effect: given a verb-noun pair, predict effect

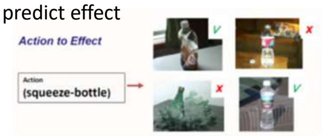

Effect to Action: given a state, predict verb-noun pair

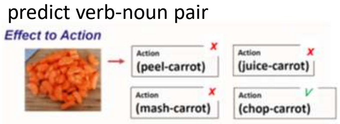

Learn from a few annotated examples to make it possible for learning through communication

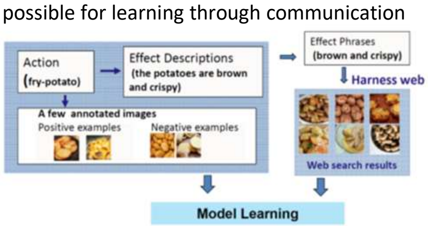

Incorporating action effect learning in teaching robots new tasks through language communication