Robot Self-Modeling

Robert Kwiatkowski and Hod Lipson
Columbia University

Specialized Functions: Lifelong Learning Machines (L2M)

Background

Data driven approaches have solved many problems in robotics. However, they require a lot of data that, in robotics, is hard to get. By using all data collected by a robot to generate a model of its own dynamics, this self-model can generate artificial data and allow classical planners or modern reinforcement learning systems to accomplish a variety of tasks. Self-modeling will allow robots to continue to improve their self-model and adapt to new tasks through their entire lifetime and truly become lifelong learning machines.

Overview

To train our self-model we collect action-state tuples corresponding to some system's action and the corresponding observed state. We then use this to learn a cause and effect relationship. This relationship allows the self-model to "look into the future" and predict what will happen if actions are taken. This can be used to plan to accomplish tasks or teach learning agents how to do tasks purely in self-simulation.

Work in Progress

After demonstrating our successes in simulation we plan to move to real world four-legged walking robots. We built the robot on the right and are in the process of collecting data and working to move our self-modeling to the real robot.

Contact:

Robert Kwiatkowski:
robert.kwiatkowski@columbia.edu
Hod Lipson:
hod.lipson@columbia.edu

Approach

Robot babbles to collect action - state pairs

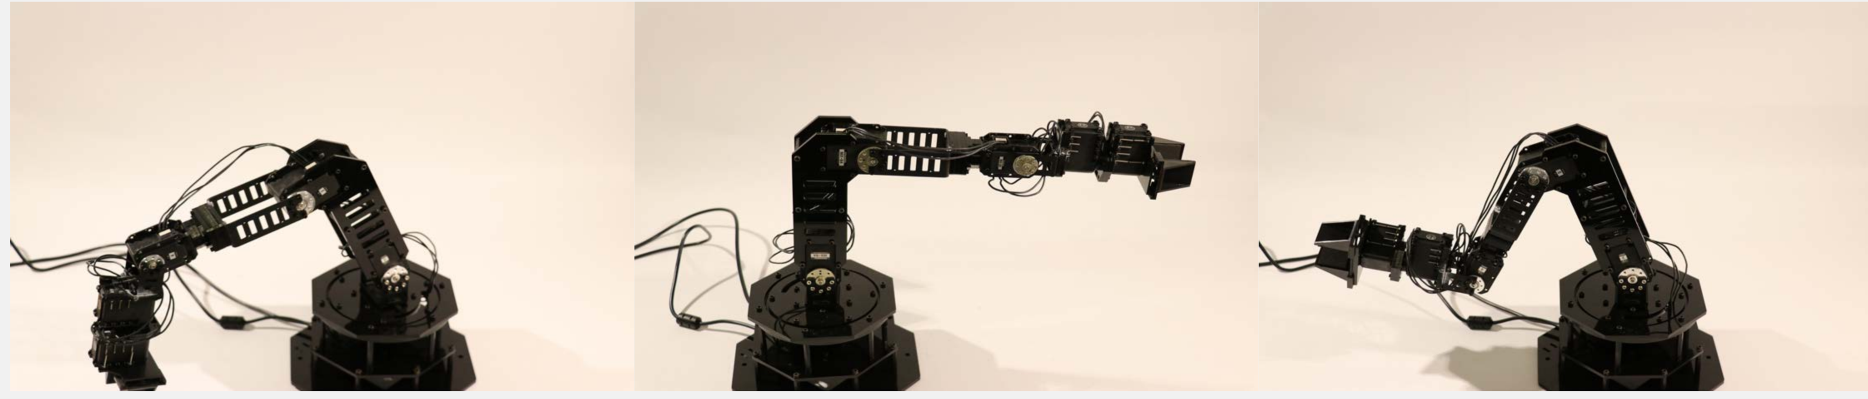

Can work with varied, complex systems

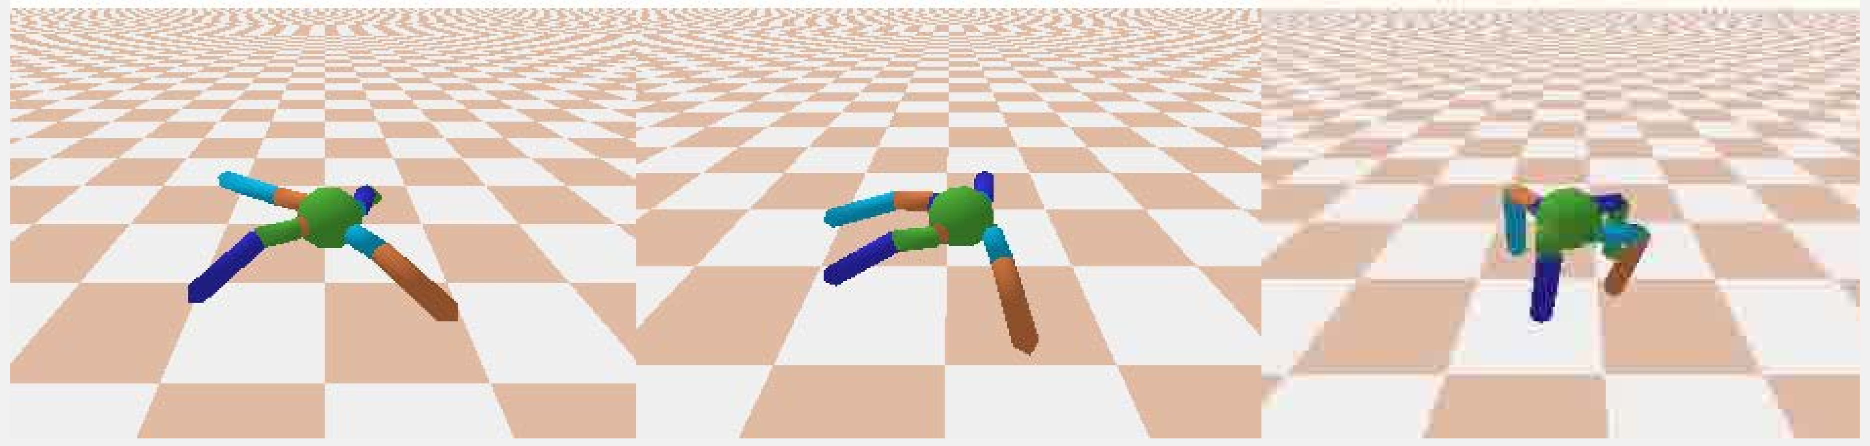

Robot learns self-model consistent with data

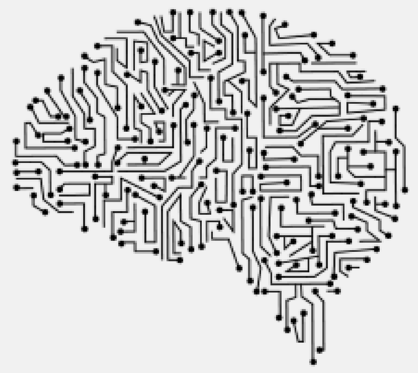

Planners or learning agents use self-model to do various tasks

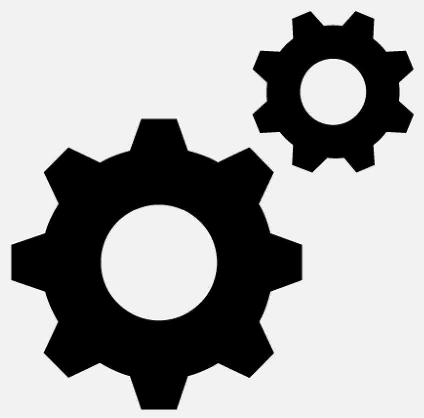

Pick-and-Place

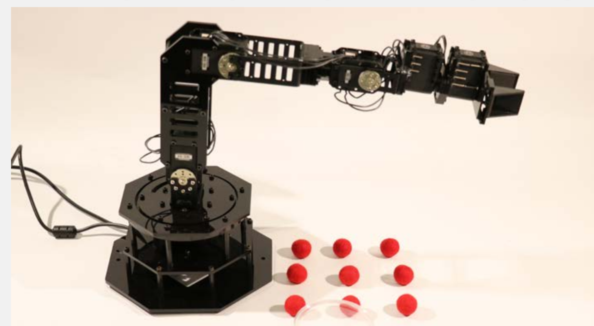

Writing

Walking

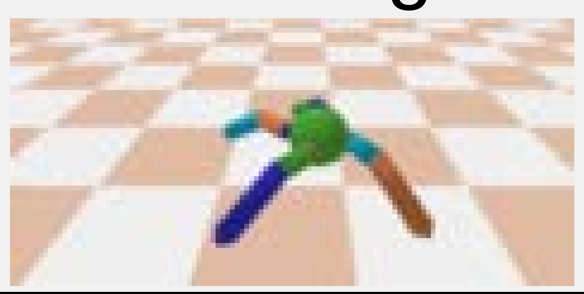

Jumping

Self-model detects changes in system, self-model is adapted with new data

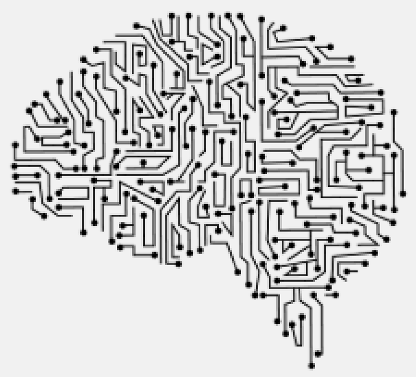

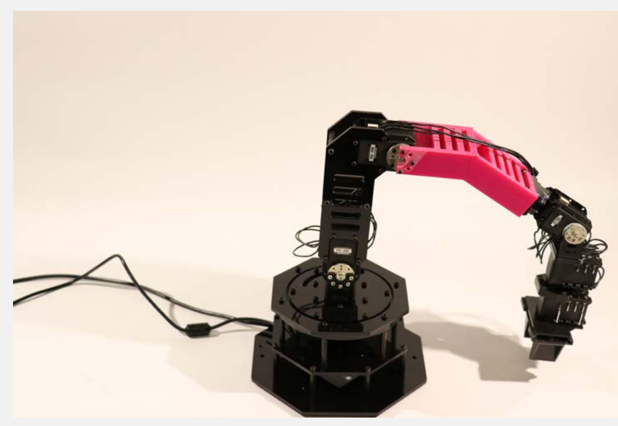

Adapted self-model can resume task

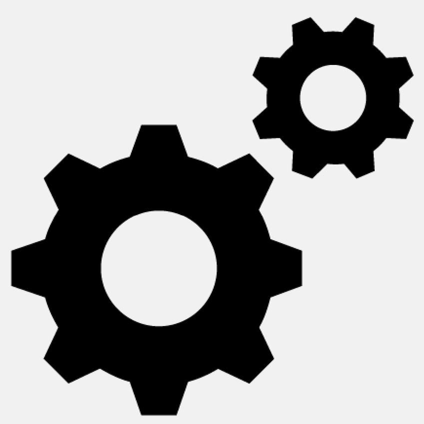

Pick-and-Place

Results

Using two different classes of systems, robotic arms and walking robots, we have successfully trained self-models. Through using these self-models we were able to program a greedy planner that was able to execute pick-and-place tasks as well as writing tasks. We then used a self-model of a four legged bullet simulation to train a reinforcement learner without exposing the RL agent to any "real" state data. The agent was able to learn how to walk and then jump without the need to collect any data more than was needed to train the self-model. These results demonstrate that a robot can use a self-model to accomplish tasks that previously needed a hand-coded simulator.

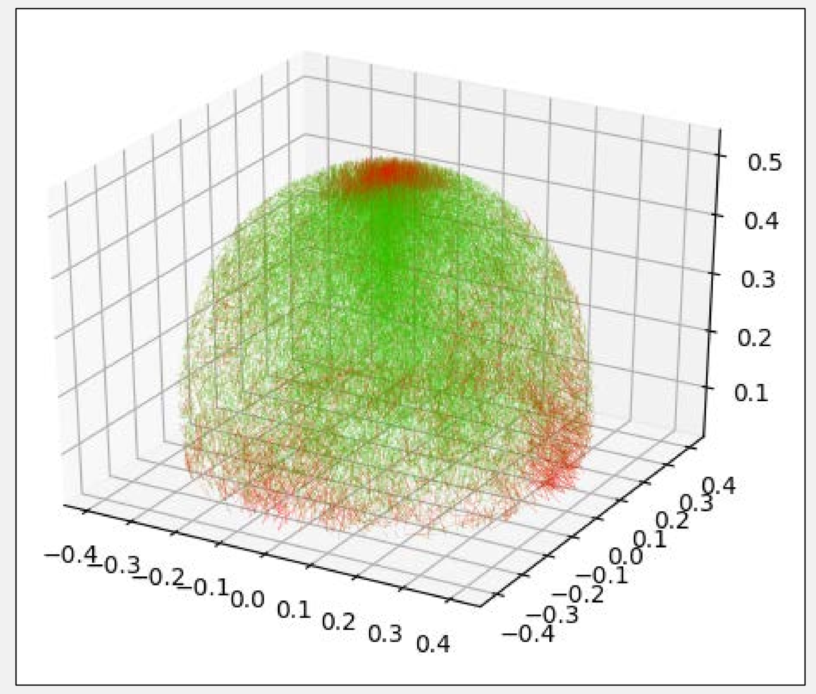

Precision of arm self-model in its reachable space. Red indicates areas of lower confidence (extreme edges).

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA), grant HR0011-18-2-0020. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

THE ELECTRONICS RESURGENCE INITIATIVE