


A new machine-learning system helps robots understand and perform certain social interactions

November 4 2021, by Adam Zewe



Credit: CC0 Public Domain

Robots can deliver food on a college campus and hit a hole in one on the golf course, but even the most sophisticated robot can't perform basic social interactions that are critical to everyday human life.

MIT researchers have now incorporated certain social interactions into a framework for robotics, enabling machines to understand what it means to help or hinder one another, and to learn to perform these social behaviors on their own. In a simulated environment, a [robot](#) watches its

companion, guesses what task it wants to accomplish, and then helps or hinders this other robot based on its own goals.

The researchers also showed that their model creates realistic and predictable social interactions. When they showed videos of these simulated robots interacting with one another to humans, the human viewers mostly agreed with the model about what type of [social behavior](#) was occurring.

Enabling robots to exhibit social skills could lead to smoother and more positive human-robot interactions. For instance, a robot in an assisted living facility could use these capabilities to help create a more caring environment for elderly individuals. The new model may also enable scientists to measure social interactions quantitatively, which could help psychologists study autism or analyze the effects of antidepressants.

"Robots will live in our world soon enough and they really need to learn how to communicate with us on human terms. They need to understand when it is time for them to help and when it is time for them to see what they can do to prevent something from happening. This is very early work and we are barely scratching the surface, but I feel like this is the first very serious attempt for understanding what it means for humans and machines to interact socially," says Boris Katz, principal research scientist and head of the InfoLab Group in the Computer Science and Artificial Intelligence Laboratory (CSAIL) and a member of the Center for Brains, Minds, and Machines (CBMM).

Joining Katz on the paper are co-lead author Ravi Tejwani, a research assistant at CSAIL; co-lead author Yen-Ling Kuo, a CSAIL Ph.D. student; Tianmin Shu, a postdoc in the Department of Brain and Cognitive Sciences; and senior author Andrei Barbu, a research scientist at CSAIL and CBMM. The research will be presented at the Conference on Robot Learning in November.

A social simulation

To study social interactions, the researchers created a simulated environment where robots pursue physical and social goals as they move around a two-dimensional grid.

A physical [goal](#) relates to the environment. For example, a robot's physical goal might be to navigate to a tree at a certain point on the grid. A social goal involves guessing what another robot is trying to do and then acting based on that estimation, like helping another robot water the tree.

The researchers use their model to specify what a robot's physical goals are, what its social goals are, and how much emphasis it should place on one over the other. The robot is rewarded for actions it takes that get it closer to accomplishing its goals. If a robot is trying to help its companion, it adjusts its reward to match that of the other robot; if it is trying to hinder, it adjusts its reward to be the opposite. The planner, an algorithm that decides which actions the robot should take, uses this continually updating reward to guide the robot to carry out a blend of physical and social goals.

"We have opened a new mathematical framework for how you model [social interaction](#) between two agents. If you are a robot, and you want to go to location X, and I am another robot and I see that you are trying to go to location X, I can cooperate by helping you get to location X faster. That might mean moving X closer to you, finding another better X, or taking whatever action you had to take at X. Our formulation allows the plan to discover the 'how'; we specify the 'what' in terms of what social interactions mean mathematically," says Tejwani.

Blending a robot's physical and social goals is important to create realistic interactions, since humans who help one another have limits to

how far they will go. For instance, a rational person likely wouldn't just hand a stranger their wallet, Barbu says.

The researchers used this mathematical framework to define three types of robots. A level 0 robot has only physical goals and cannot reason socially. A level 1 robot has physical and social goals but assumes all other robots only have physical goals. Level 1 robots can take actions based on the physical goals of other robots, like helping and hindering. A level 2 robot assumes other robots have social and physical goals; these robots can take more sophisticated actions like joining in to help together.

Evaluating the model

To see how their model compared to human perspectives about social interactions, they created 98 different scenarios with robots at levels 0, 1, and 2. Twelve humans watched 196 video clips of the robots interacting, and then were asked to estimate the physical and social goals of those robots.

In most instances, their model agreed with what the humans thought about the social interactions that were occurring in each frame.

"We have this long-term interest, both to build computational models for robots, but also to dig deeper into the human aspects of this. We want to find out what features from these videos humans are using to understand social interactions. Can we make an objective test for your ability to recognize social interactions? Maybe there is a way to teach people to recognize these social interactions and improve their abilities. We are a long way from this, but even just being able to measure social interactions effectively is a big step forward," Barbu says.

Toward greater sophistication

The researchers are working on developing a system with 3D agents in an environment that allows many more types of interactions, such as the manipulation of household objects. They are also planning to modify their model to include environments where actions can fail.

The researchers also want to incorporate a neural network-based robot planner into the [model](#), which learns from experience and performs faster. Finally, they hope to run an experiment to collect data about the features humans use to determine if two robots are engaging in a social interaction.

"Hopefully, we will have a benchmark that allows all researchers to work on these social interactions and inspire the kinds of science and engineering advances we've seen in other areas such as object and action recognition," Barbu says.

More information: Ravi Tejjwani et al, Social Interactions as Recursive MDPs (2021). Available as a PDF at openreview.net/pdf?id=3HZLte8gMYS

Provided by Massachusetts Institute of Technology

Citation: A new machine-learning system helps robots understand and perform certain social interactions (2021, November 4) retrieved 19 April 2024 from <https://techxplore.com/news/2021-11-machine-learning-robots-social-interactions.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.