

Model helps robots navigate more like humans do

October 4 2018, by Rob Matheson



MIT researchers have devised a way to help robots navigate environments more like humans do. Credit: Massachusetts Institute of Technology

When moving through a crowd to reach some end goal, humans can usually navigate the space safely without thinking too much. They can

learn from the behavior of others and note any obstacles to avoid. Robots, on the other hand, struggle with such navigational concepts.

MIT researchers have now devised a way to help robots navigate environments more like humans do. Their novel motion-planning model lets robots determine how to reach a goal by exploring the environment, observing other agents, and exploiting what they've learned before in similar situations. A paper describing the model was presented at this week's IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).

Popular motion-planning algorithms will create a tree of possible decisions that branches out until it finds good paths for navigation. A [robot](#) that needs to navigate a room to reach a door, for instance, will create a step-by-step search tree of possible movements and then execute the best path to the door, considering various constraints. One drawback, however, is these algorithms rarely learn: Robots can't leverage information about how they or other agents acted previously in similar environments.

"Just like when playing chess, these decisions branch out until [the robots] find a good way to navigate. But unlike chess players, [the robots] explore what the future looks like without learning much about their environment and other agents," says co-author Andrei Barbu, a researcher at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and the Center for Brains, Minds, and Machines (CBMM) within MIT's McGovern Institute. "The thousandth time they go through the same crowd is as complicated as the first time. They're always exploring, rarely observing, and never using what's happened in the past."

The researchers developed a model that combines a planning algorithm with a [neural network](#) that learns to recognize paths that could lead to the

best outcome, and uses that knowledge to guide the robot's movement in an environment.

In their paper, "Deep sequential models for sampling-based planning," the researchers demonstrate the advantages of their model in two settings: navigating through challenging rooms with traps and narrow passages, and navigating areas while avoiding collisions with other agents. A promising real-world application is helping [autonomous cars](#) navigate intersections, where they have to quickly evaluate what others will do before merging into traffic. The researchers are currently pursuing such applications through the Toyota-CSAIL Joint Research Center.

"When humans interact with the world, we see an object we've interacted with before, or are in some location we've been to before, so we know how we're going to act," says Yen-Ling Kuo, a Ph.D. in CSAIL and first author on the paper. "The idea behind this work is to add to the search space a machine-learning model that knows from past experience how to make planning more efficient."

Boris Katz, a principal research scientist and head of the InfoLab Group at CSAIL, is also a co-author on the paper.

Trading off exploration and exploitation

Traditional motion planners explore an environment by rapidly expanding a tree of decisions that eventually blankets an entire space. The robot then looks at the tree to find a way to reach the goal, such as a door. The researchers' model, however, offers "a tradeoff between exploring the world and exploiting past knowledge," Kuo says.

The learning process starts with a few examples. A robot using the model is trained on a few ways to navigate similar environments. The

neural network learns what makes these examples succeed by interpreting the environment around the robot, such as the shape of the walls, the actions of other agents, and features of the goals. In short, the model "learns that when you're stuck in an environment, and you see a doorway, it's probably a good idea to go through the door to get out," Barbu says.

The model combines the exploration behavior from earlier methods with this learned information. The underlying planner, called RRT*, was developed by MIT professors Sertac Karaman and Emilio Frazzoli. (It's a variant of a widely used motion-planning algorithm known as Rapidly-exploring Random Trees, or RRT.) The planner creates a search tree while the neural network mirrors each step and makes probabilistic predictions about where the robot should go next. When the network makes a prediction with high confidence, based on learned information, it guides the robot on a new path. If the network doesn't have high confidence, it lets the robot explore the environment instead, like a traditional planner.

For example, the researchers demonstrated the model in a simulation known as a "bug trap," where a 2-D robot must escape from an inner chamber through a central narrow channel and reach a location in a surrounding larger room. Blind allies on either side of the channel can get robots stuck. In this simulation, the robot was trained on a few examples of how to escape different bug traps. When faced with a new trap, it recognizes features of the trap, escapes, and continues to search for its goal in the larger room. The neural network helps the robot find the exit to the trap, identify the dead ends, and gives the robot a sense of its surroundings so it can quickly find the goal.

Results in the paper are based on the chances that a path is found after some time, total length of the path that reached a given goal, and how consistent the paths were. In both simulations, the researchers' model

more quickly plotted far shorter and consistent paths than a traditional planner.

Working with multiple agents

In one other experiment, the researchers trained and tested the model in navigating environments with multiple moving agents, which is a useful test for autonomous cars, especially navigating intersections and roundabouts. In the simulation, several agents are circling an obstacle. A robot agent must successfully navigate around the other agents, avoid collisions, and reach a goal location, such as an exit on a roundabout.

"Situations like roundabouts are hard, because they require reasoning about how others will respond to your actions, how you will then respond to theirs, what they will do next, and so on," Barbu says. "You eventually discover your first action was wrong, because later on it will lead to a likely accident. This problem gets exponentially worse the more cars you have to contend with."

Results indicate that the researchers' model can capture enough information about the future behavior of the other agents (cars) to cut off the process early, while still making good decisions in navigation. This makes planning more efficient. Moreover, they only needed to train the model on a few examples of roundabouts with only a few cars. "The plans the robots make take into account what the other cars are going to do, as any human would," Barbu says.

Going through intersections or roundabouts is one of the most challenging scenarios facing autonomous cars. This work might one day let cars learn how humans behave and how to adapt to drivers in different environments, according to the researchers. This is the focus of the Toyota-CSAIL Joint Research Center work.

"Not everybody behaves the same way, but people are very stereotypical. There are people who are shy, people who are aggressive. The [model](#) recognizes that quickly and that's why it can plan efficiently," Barbu says.

More recently, the researchers have been applying this work to robots with manipulators that face similarly daunting challenges when reaching for objects in ever-changing environments.

More information: Deep sequential models for sampling-based planning. arxiv.org/abs/1810.00804

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Model helps robots navigate more like humans do (2018, October 4) retrieved 28 May 2024 from <https://techxplore.com/news/2018-10-robots-humans.html>

<p>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.</p>
--