Researchers use novel approach to teach robot to navigate over obstacles

May 19 2023

From left to right, Simar Kareer, Joanne Truong, and Naoki Yokoyam work together on developing a quadrupedal robot that can navigate over obstacles. Credit: Kevin Beasley/College of Computing, Georgia Institute of Technology

Quadrupedal robots may be able to step directly over obstacles in their

paths thanks to the efforts of a trio of Georgia Tech Ph.D. students.

When it comes to robotic locomotion and [navigation](), Naoki Yokoyama says most four-legged robots are trained to regain their footing if an obstacle causes them to stumble. Working toward a larger effort to develop a housekeeping [robot](), Yokoyama and his collaborators—Simar Kareer and Joanne Truong—set out to train their robot to walk over clutter it might encounter in a home.

"The main motivation of the project is getting low-level control over the legs of the robot that also incorporates visual input," said Yokoyama, a Ph.D. student within the School of Electrical and Computer Engineering. "We envisioned a controller that could be deployed in an indoor setting with a lot of clutter, such as shoes or toys on the ground of a messy home. Whereas blind locomotive controllers tend to be more reactive—if they step on something, they'll make sure they don't fall over—we wanted ours to use [visual input]() to avoid stepping on the obstacle altogether."

To achieve their goal, the researchers took a novel training approach of fusing a high-level visual navigation policy with a visual locomotion policy.

In a paper advised by Interactive Computing Associate Professor Dhruv Batra and Assistant Professor Sehoon Ha, Kareer, Yokoyama, and Truong show that their two-policy approach successfully simulates robotic navigation over obstacles.

They call their approach ViNL (Visual Navigation and Locomotion), and so far, it has guided robots through simulated novel cluttered environments with a 72.6% success rate. The team will present its paper, ViNL: Visual Navigation and Locomotion Over Obstacles, at the [IEEE International Conference on Robotics and Automation](), which is being

held May 29–June 2 in London.

Both policies are model-free—the robot learns on its own simulation and doesn't mimic any pre-existing behavioral patterns—and can be combined without any additional co-training.

"This work uniquely combines separate locomotion and navigation policies in a zero-shot manner," said Kareer, who along with Truong is a Ph.D. student within the School of Interactive Computing. "If we come up with an improved navigation policy, we can just take that, do no extra work, and deploy that to our robot. That's a scalable approach. You can plug and play these things together with very little fine-tuning. That's powerful."

The visual navigation policy teaches the robot through goal-achieving motivation. It gives the robot an objective of navigating from one place to another while avoiding any obstacles. The robot receives a score based on how successfully it completes its task. If it stumbles over an obstacle, it is penalized.

"We gave it an environment that had very few obstacles, and then slightly more and slightly more," Kareer said. "This gradual approach is helpful to its learning. When you just toss it into an environment with a million obstacles, it fails a lot. But if you show it one or two obstacles and say, 'try to learn these,' it's much more stable."

The locomotion policy teaches the robot how to use its limbs to step over an object, including how high it should lift its legs.

Because a real-world quadruped will only be able to see what its front camera sees, obstacles will disappear from its view as it gets closer to them. The team accounted for this by incorporating memory and spatial awareness into their network architecture to teach the robot exactly when

and where to step over the obstacle.

"The robot has a rich understanding of where its entire limb is relative to the obstacles," Kareer said. "When you see it walking over obstacles, it's not just deciding to put its foot down on spots where there are no obstacles. It's remembering where all the obstacles are relative to its body and keeping its limbs out of the way until it's passed over them."

And if an obstacle is too tall to step over, the robot can also choose to go around it.

"We saw that it was very good at navigating, and even in cases where it might take a wrong turn, it knows that it can backtrack and go back where it came from," Truong said.

Finally, the group taught the robot specifically what types of objects it should be looking to step over in a house, such as toys, and ones that it should go around, such as a chair. This also helps the robot to know how high it will need to lift its legs.

"What's important for navigation is to be able to have the experience of navigating in real-world houses, so we train our navigation policy with photo-realistic scans of apartments," Truong said. "We used scans of over 1,000 apartments for training and evaluated the robot in scenarios it had never seen before. We zero-shot deploy it into a new environment, so you can take a new robot, put it in a new house, and it will be able to do this as well."

The researchers agree their paper is multi-faceted and has numerous implications that fall outside its focus but are nonetheless important. Their work could lead to robots navigating openly in the outdoors, selectively picking paths based on the user's preference to avoid muddy ground or rocky terrain.

"Normally, it matters much less how you get from Point A to Point B," Truong said. "You just need to know that Point B is valid. With overcoming obstacles, not only do Point A and Point B need to be valid, how you get from Point A to Point B also matters."

More information: ViNL: Visual Navigation and Locomotion Over Obstacles: www.joannetruong.com/projects/vinl.html

Provided by Georgia Institute of Technology