AI approach yields 'athletically intelligent' robotic dog

October 5 2023, by Andrew Myers

A timelapse photo showing how the dog-like robot is able to navigate a variety of obstacles autonomously. Credit: Shanghai Qi Zhi Institute/Stanford University

Someday, when quakes, fires, and floods strike, the first responders might be packs of robotic rescue dogs rushing in to help stranded souls. These battery-powered quadrupeds would use computer vision to size up obstacles and employ doglike agility skills to get past them.

Toward that noble goal, AI researchers at Stanford University and Shanghai Qi Zhi Institute say they have developed a new vision-based algorithm that helps robodogs scale high objects, leap across gaps, crawl under thresholds, and squeeze through crevices—and then bolt to the next challenge. The algorithm represents the brains of the robodog.

"The autonomy and range of complex skills that our quadruped robot

learned is quite impressive," said Chelsea Finn, assistant professor of computer science and senior author of a new peer-reviewed paper announcing the teams' approach to the world, which will be presented at the upcoming Conference on Robot Learning held Nov. 6–9 in Atlanta. "And we have created it using low-cost, off-the-shelf robots—actually, two different off-the-shelf robots."

The key advance, the authors say, is that their robodog is autonomous—that is, it is able to size up physical challenges and imagine, then execute, a broad range of agility skills based simply on the obstacles it sees before it.

"What we're doing is combining both perception and control, using images from a depth camera mounted on the robot and machine learning to process all those inputs and move the legs in order to get over, under, and around obstacles," said Zipeng Fu, a doctoral candidate in Finn's lab and first author of the study, along with Ziwen Zhuang of Shanghai Qi Zhi Institute.

Simplifying to optimize

Theirs is not the first robodog to demonstrate such agility—a class of athletics known as "parkour"—but it is first to combine self-sufficiency with a broad array of skills.

"Our robots have both vision and autonomy—the athletic intelligence to size up a challenge and to self-select and execute parkour skills based on the demands of the moment," Fu said.

Existing learning methods are often based on complex reward systems that must be fine-tuned to specific physical obstacles. Accordingly, they don't scale well to new or unfamiliar environments. Other related approaches learn using real-world data to imitate agility skills of other

animals. These robodogs lack a broad skill set and don't have the new robodogs' vision capabilities. Both existing methods are also computationally "laggy"—in other words, slow.

This is the first open-source application to accomplish these goals with a simple reward system using no real-world reference data, the authors write in the study.

To succeed, they first synthesized and honed the algorithm using a computer model, then transferred it to two real-world robodogs. Next, in a process called reinforcement learning, the robots attempted to move forward in any way they saw fit and got rewarded based on how well they did. This is how the algorithm eventually learns the best way to approach a new challenge.

In practice, most existing [reinforcement learning](link) reward systems involve too many variables to be effective, slowing computational performance. This is what makes the streamlined reward process for robodog parkour exceptional, if also surprisingly straightforward.

"It's actually fairly simple," Finn said. "We based it mostly on how far forward the robot is moving and the amount of effort it has applied to do it. Eventually, the robot learns more complex motor skills that allow it to get ahead."

Real-world tests

The team then performed extensive experiments using real-world robodogs to demonstrate their new agility approach in especially challenging environments using only those robodogs' off-the-shelf computers, visual sensors, and power systems.

In raw numbers, the new-and-improved robodogs were able to climb

obstacles more than one-and-a-half times their height, leap gaps greater than one-and-a-half times their length, crawl beneath barriers three-quarters of their height, and tilt themselves in order to squeeze through a slit thinner than their width.

Next up, the team hopes to leverage advances in 3D vision and graphics to add real-world data to its simulated environments to bring a new level of real-world autonomy to their algorithm.

  More information: Robot Parkour Learning. openreview.net/forum?id=uo937r5eTE

Provided by Stanford University