


Control system enables four-legged robots to jump across uneven terrain in real time

October 21 2021, by Adam Zewe



MIT researchers have developed a system that improves the speed and agility of legged robots as they jump across gaps in the terrain. Credit: Massachusetts Institute of Technology

A loping cheetah dashes across a rolling field, bounding over sudden

gaps in the rugged terrain. The movement may look effortless, but getting a robot to move this way is an altogether different prospect.

In recent years, four-legged robots inspired by the movement of cheetahs and other animals have made great leaps forward, yet they still lag behind their mammalian counterparts when it comes to traveling across a landscape with rapid elevation changes.

"In those settings, you need to use vision in order to avoid failure. For example, stepping in a gap is difficult to avoid if you can't see it. Although there are some existing methods for incorporating vision into legged locomotion, most of them aren't really suitable for use with emerging agile robotic systems," says Gabriel Margolis, a Ph.D. student in the lab of Pulkit Agrawal, professor in the Computer Science and Artificial Intelligence Laboratory (CSAIL) at MIT.

Now, Margolis and his collaborators have developed a system that improves the speed and agility of legged robots as they jump across gaps in the terrain. The novel control system is split into two parts—one that processes real-time input from a [video camera](#) mounted on the front of the [robot](#) and another that translates that information into instructions for how the robot should move its body. The researchers tested their system on the MIT mini cheetah, a powerful, agile robot built in the lab of Sangbae Kim, professor of mechanical engineering.

Unlike other methods for controlling a four-legged robot, this two-part system does not require the terrain to be mapped in advance, so the robot can go anywhere. In the future, this could enable robots to charge off into the woods on an emergency response mission or climb a flight of stairs to deliver medication to an elderly shut-in.

Margolis wrote the paper with senior author Pulkit Agrawal, who heads the Improbable AI lab at MIT and is the Steven G. and Renee Finn

Career Development Assistant Professor in the Department of Electrical Engineering and Computer Science; Professor Sangbae Kim in the Department of Mechanical Engineering at MIT; and fellow graduate students Tao Chen and Xiang Fu at MIT. Other co-authors include Kartik Paigwar, a graduate student at Arizona State University; and Donghyun Kim, an assistant professor at the University of Massachusetts at Amherst. The work will be presented next month at the Conference on Robot Learning.

It's all under control

The use of two separate controllers working together makes this system especially innovative.

A [controller](#) is an algorithm that will convert the robot's state into a set of actions for it to follow. Many blind controllers—those that do not incorporate vision—are robust and effective but only enable robots to walk over continuous terrain.

Vision is such a complex sensory input to process that these algorithms are unable to handle it efficiently. Systems that do incorporate vision usually rely on a "heightmap" of the terrain, which must be either preconstructed or generated on the fly, a process that is typically slow and prone to failure if the heightmap is incorrect.

To develop their system, the researchers took the best elements from these robust, blind controllers and combined them with a separate module that handles vision in real-time.

The robot's camera captures depth images of the upcoming terrain, which are fed to a high-level controller along with information about the state of the robot's body (joint angles, body orientation, etc.). The high-level controller is a [neural network](#) that "learns" from experience.

That neural network outputs a target trajectory, which the second controller uses to come up with torques for each of the robot's 12 joints. This low-level controller is not a neural network and instead relies on a set of concise, physical equations that describe the robot's motion.

"The hierarchy, including the use of this low-level controller, enables us to constrain the robot's behavior so it is more well-behaved. With this low-level controller, we are using well-specified models that we can impose constraints on, which isn't usually possible in a learning-based network," Margolis says.

Teaching the network

The researchers used the trial-and-error method known as reinforcement learning to train the high-level controller. They conducted simulations of the robot running across hundreds of different discontinuous terrains and rewarded it for successful crossings.

Over time, the algorithm learned which actions maximized the reward.

Then they built a physical, gapped terrain with a set of wooden planks and put their control scheme to the test using the mini cheetah.

"It was definitely fun to work with a robot that was designed in-house at MIT by some of our collaborators. The mini cheetah is a great platform because it is modular and made mostly from parts that you can order online, so if we wanted a new battery or camera, it was just a simple matter of ordering it from a regular supplier and, with a little bit of help from Sangbae's lab, installing it," Margolis says.

Estimating the robot's state proved to be a challenge in some cases. Unlike in simulation, real-world sensors encounter noise that can accumulate and affect the outcome. So, for some experiments that

involved high-precision foot placement, the researchers used a motion capture system to measure the robot's true position.

Their system outperformed others that only use one controller, and the mini cheetah successfully crossed 90 percent of the terrains.

"One novelty of our system is that it does adjust the robot's gait. If a human were trying to leap across a really wide gap, they might start by running really fast to build up speed and then they might put both feet together to have a really powerful leap across the gap. In the same way, our robot can adjust the timings and duration of its foot contacts to better traverse the terrain," Margolis says.

Leaping out of the lab

While the researchers were able to demonstrate that their control scheme works in a laboratory, they still have a long way to go before they can deploy the system in the real world, Margolis says.

In the future, they hope to mount a more powerful computer to the robot so it can do all its computation on board. They also want to improve the robot's state estimator to eliminate the need for the motion capture system. In addition, they'd like to improve the low-level controller so it can exploit the robot's full range of motion, and enhance the high-level controller so it works well in different lighting conditions.

"It is remarkable to witness the flexibility of machine learning techniques capable of bypassing carefully designed intermediate processes (e.g. state estimation and trajectory planning) that centuries-old model-based techniques have relied on," Kim says. "I am excited about the future of mobile robots with more robust vision processing trained specifically for locomotion."

More information: Learning to Jump from Pixels.
openreview.net/forum?id=R4E8wTUtxdl

*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: Control system enables four-legged robots to jump across uneven terrain in real time
(2021, October 21) retrieved 29 April 2024 from <https://techxplore.com/news/2021-10-enables-four-legged-robots-uneven-terrain.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.