


3 Questions: How the mini cheetah robot learns to run

March 17 2022, by Rachel Gordon



MIT's mini cheetah, using a model-free reinforcement learning system, broke the record for the fastest run recorded. Credit: MIT CSAIL

It's been roughly 23 years since one of the first robotic animals trotted on the scene, defying classical notions of our cuddly four-legged friends.

Since then, a barrage of the walking, dancing, and door-opening machines have commanded their presence, a sleek mixture of batteries, sensors, metal, and motors. Missing from the list of cardio activities was one both loved and loathed by humans (depending on whom you ask), and which proved slightly trickier for the bots: learning to run.

Researchers from MIT's Improbable AI Lab, part of the Computer Science and Artificial Intelligence Laboratory (CSAIL) and directed by MIT Assistant Professor Pulkit Agrawal, as well as the Institute of AI and Fundamental Interactions (IAIFI) have been working on fast-paced strides for a robotic mini cheetah—and their model-free reinforcement learning system broke the record for the fastest run recorded. Here, MIT Ph.D. student Gabriel Margolis and IAIFI postdoc Ge Yang discuss just how fast the cheetah can run.

Q: We've seen videos of robots running before. Why is running harder than walking?

A: Achieving fast running requires pushing the hardware to its limits, for example by operating near the maximum torque output of motors. In such conditions, the robot dynamics are hard to analytically model. The robot needs to respond quickly to changes in the environment, such as the moment it encounters ice while running on grass. If the robot is walking, it is moving slowly and the presence of snow is not typically an issue. Imagine if you were walking slowly, but carefully: you can traverse almost any terrain. Today's robots face an analogous problem. The problem is that moving on all terrains as if you were walking on ice is very inefficient, but is common among today's robots. Humans run fast on grass and slow down on ice—we adapt. Giving robots a similar capability to adapt requires quick identification of terrain changes and quickly adapting to prevent the robot from falling over. In summary, because it's impractical to build analytical (human-designed) models of all possible terrains in advance, and the robot's dynamics become more

complex at high-velocities, high-speed running is more challenging than walking.

Q: Previous agile running controllers for the MIT Cheetah 3 and mini cheetah, as well as for Boston Dynamics' robots, are "analytically designed," relying on human engineers to analyze the physics of locomotion, formulate efficient abstractions, and implement a specialized hierarchy of controllers to make the robot balance and run. You use a "learn-by-experience model" for running instead of programming it. Why?

A: Programming how a robot should act in every possible situation is simply very hard. The process is tedious, because if a robot were to fail on a particular terrain, a human engineer would need to identify the cause of failure and manually adapt the robot controller, and this process can require substantial human time. Learning by trial and error removes the need for a human to specify precisely how the robot should behave in every situation. This would work if: (1) the robot can experience an extremely wide range of terrains; and (2) the robot can automatically improve its behavior with experience.

Thanks to modern simulation tools, our robot can accumulate 100 days' worth of experience on diverse terrains in just three hours of actual time. We developed an approach by which the robot's behavior improves from simulated experience, and our approach critically also enables successful deployment of those learned behaviors in the real world. The intuition behind why the robot's running skills work well in the real world is: Of all the environments it sees in this simulator, some will teach the robot skills that are useful in the real world. When operating in the real world, our controller identifies and executes the relevant skills in real-time.

Q: Can this approach be scaled beyond the mini cheetah? What excites you about its future applications?

A: At the heart of [artificial intelligence research](#) is the trade-off between what the human needs to build in (nature) and what the machine can learn on its own (nurture). The traditional paradigm in robotics is that humans tell the robot both what task to do and how to do it. The problem is that such a framework is not scalable, because it would take immense human engineering effort to manually program a robot with the skills to operate in many diverse environments. A more practical way to build a robot with many diverse skills is to tell the robot what to do and let it figure out the how. Our system is an example of this. In our lab, we've begun to apply this paradigm to other robotic systems, including hands that can pick up and manipulate many different objects.

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: 3 Questions: How the mini cheetah robot learns to run (2022, March 17) retrieved 20 April 2024 from <https://techxplore.com/news/2022-03-mini-cheetah-robot.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.