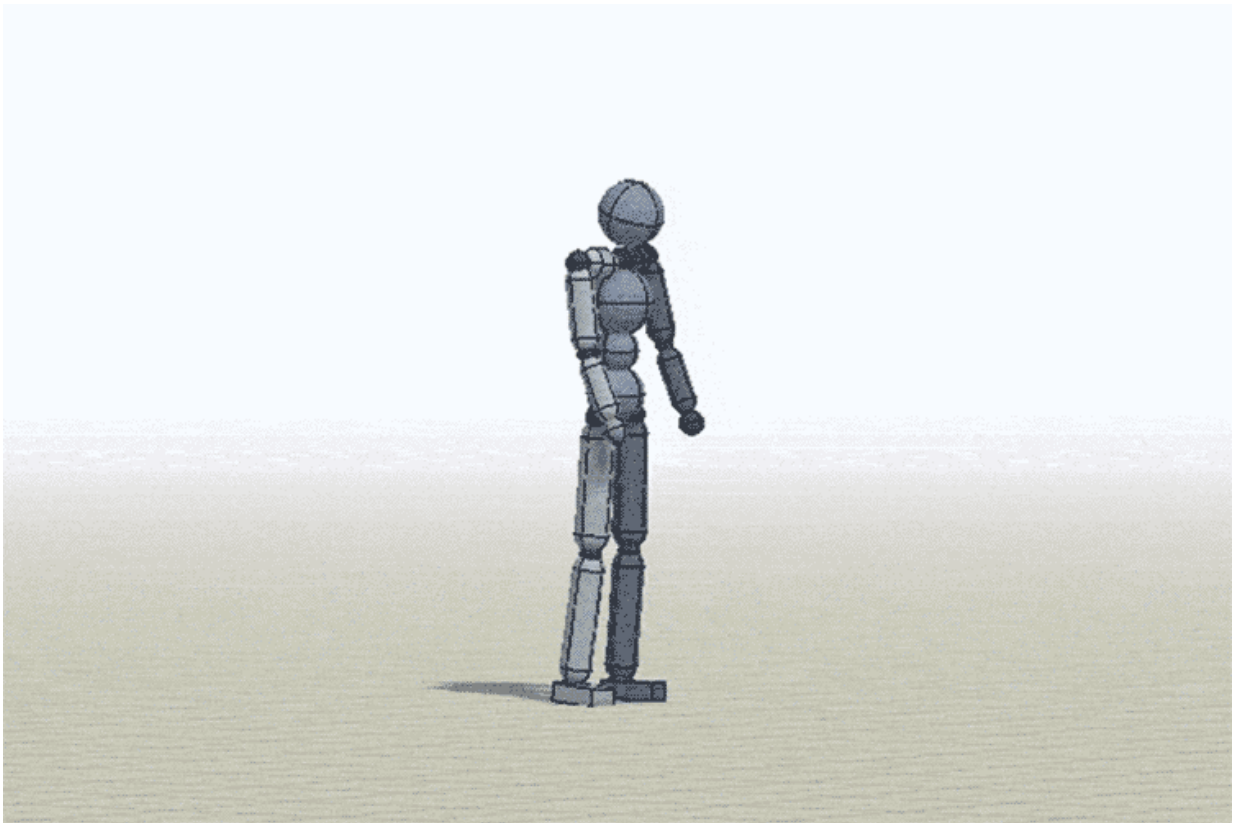# Making computer animation more agile, acrobatic—and realistic

April 10 2018

UC Berkeley computer scientists developed an algorithm that uses reinforcement learning to generate realistic simulations of human and animal motion, such as this real-time backflip. The same algorithm works for 25 acrobatic and dance tricks, with one month of learning required per skill. Credit: Jason Peng, UC Berkeley

It's still easy to tell computer-simulated motions from the real thing - on the big screen or in video games, simulated humans and animals often move clumsily, without the rhythm and fluidity of their real-world counterparts.

But that's changing. University of California, Berkeley researchers have now made a major advance in realistic [computer animation](#), using deep reinforcement learning to recreate natural motions, even for acrobatic feats like break dancing and martial arts. The simulated characters can also respond naturally to changes in the environment, such as recovering from tripping or being pelted by projectiles.

"This is actually a pretty big leap from what has been done with deep learning and animation. In the past, a lot of work has gone into simulating natural motions, but these physics-based methods tend to be very specialized; they're not general methods that can handle a large variety of skills," said UC Berkeley graduate student Xue Bin "Jason" Peng. Each activity or task typically requires its own custom-designed controller.

"We developed more capable agents that behave in a natural manner," he said. "If you compare our results to motion-capture recorded from humans, we are getting to the point where it is pretty difficult to distinguish the two, to tell what is simulation and what is real. We're moving toward a virtual stuntman."

The work could also inspire the development of more dynamic motor skills for robots.

A paper describing the development has been conditionally accepted for presentation at the 2018 SIGGRAPH conference in August in Vancouver, Canada, and was posted online April 10. Peng's colleagues in the Department of Electrical Engineering and Computer Sciences are

professor Pieter Abbeel and assistant professor Sergey Levine, along with Michiel van de Panne of the University of British Columbia.

Mocap for DeepMimic

Traditional techniques in animation typically require designing custom controllers by hand for every skill: one controller for walking, for example, and another for running, flips and other movements. These hand-designed controllers can look pretty good, Peng said.

Alternatively, deep reinforcement learning methods, such as GAIL, can simulate a variety of different skills using a single general algorithm, but their results often look very unnatural.

"The advantage of our work," Peng said, "is that we can get the best of both worlds. We have a single algorithm that can learn a variety of different skills, and produce motions that rival if not surpass the state of the art in animation with handcrafted controllers."

To achieve this, Peng obtained reference data from motion-capture (mocap) clips demonstrating more than 25 different acrobatic feats, such as backflips, cartwheels, kip-ups and vaults, as well as simple running, throwing and jumping. After providing the mocap data to the computer, the team then allowed the system - dubbed DeepMimic - to "practice" each skill for about a month of simulated time, a bit longer than a human might take to learn the same skill.

The computer practiced 24/7, going through millions of trials to learn how to realistically simulate each skill. It learned through trial and error: comparing its performance after each trial to the mocap data, and tweaking its behavior to more closely match the human motion.

"The machine is learning these skills completely from scratch, before it

even knows how to walk or run, so a month might not be too unreasonable," he said.

The key was allowing the machine to learn in ways that humans don't. For example, a backflip involves so many individual body movements that a machine might keep falling and never get past the first few steps. Instead, the algorithm starts learning at various stages of the backflip - including in mid-air - so as to learn each stage of the motion separately and then stitch them together.

Surprisingly, once trained, the simulated characters are able to deal with and recover from never-before-seen conditions: running over irregular terrain and doing spin-kicks while being pelted by projectiles.

"The recoveries come for free from the learning process," Peng said.

And the same simple method worked for all of the more than 25 skills.

"When we first started, we thought we would try something simple, as a baseline for later methods, not expecting that it was going to work. But the very simple method actually works really well. This shows that a simple approach can actually learn a very rich repertoire of highly dynamic and acrobatic skills."

Provided by University of California - Berkeley