ARTIFICIAL INTELLIGENCE

Robots on the run

After decades of clumsiness, robots are finally learning to walk, run and grasp with grace. Such progress spells the beginning of an age of physically adept artificial intelligence.

HOD LIPSON

Young animals gallop across fields, climb trees and immediately find their feet with enviable grace after they fall[1]. And like our primate cousins, humans can deploy opposable thumbs and fine motor skills to complete tasks such as effortlessly peeling a clementine or feeling for the correct key in a dark hallway. Although walking and grasping are easy for many living things, robots have been notoriously poor at gaited locomotion and manual dexterity. Until now.

Writing in *Science Robotics*, Hwangbo *et al.*[2] report intriguing evidence that a data-driven approach to designing robotic software could overcome a long-standing challenge in robotics and artificial-intelligence research called the simulation–reality gap. For decades, roboticists have guided the limbs of robots using software that is built on a foundation of predictive, mathematical models, known as classical control theory. However, this method has proved ineffective when applied to the seemingly simple problem of guiding robotic limbs through the tasks of walking, climbing and grasping objects of various shapes.

A robot typically begins its life in simulation. When its guiding software performs well in the virtual world, that software is placed in a robotic body and then sent into the physical world. There, the robot will inevitably encounter limitless, and difficult to predict, irregularities in the environment. Examples of such issues include surface friction, structural flexibility, vibration, sensor delays and poorly timed actuators — devices that convert energy into movement. Unfortunately, these combined nuisances are impossible to describe fully, in advance, using mathematics. As a result, even a robot that performs beautifully in simulation will stumble and fall after a few encounters with seemingly minor physical obstacles.

Hwangbo *et al.* have demonstrated a way of closing this performance gap by blending classical control theory with machine-learning techniques. The team began by designing a conventional mathematical model of a medium-dog-sized quadrupedal robot called ANYmal (Fig. 1). Next, they collected data from the actuators that guide the movements of the robot's limbs. They fed this information into several machine-learning systems known as neural networks to build a second model — one that could automatically predict the idiosyncratic movements of the AMYmal robot's limbs. Finally, the team inserted the trained neural networks into its first model and ran the hybrid model in simulation on a standard desktop computer.

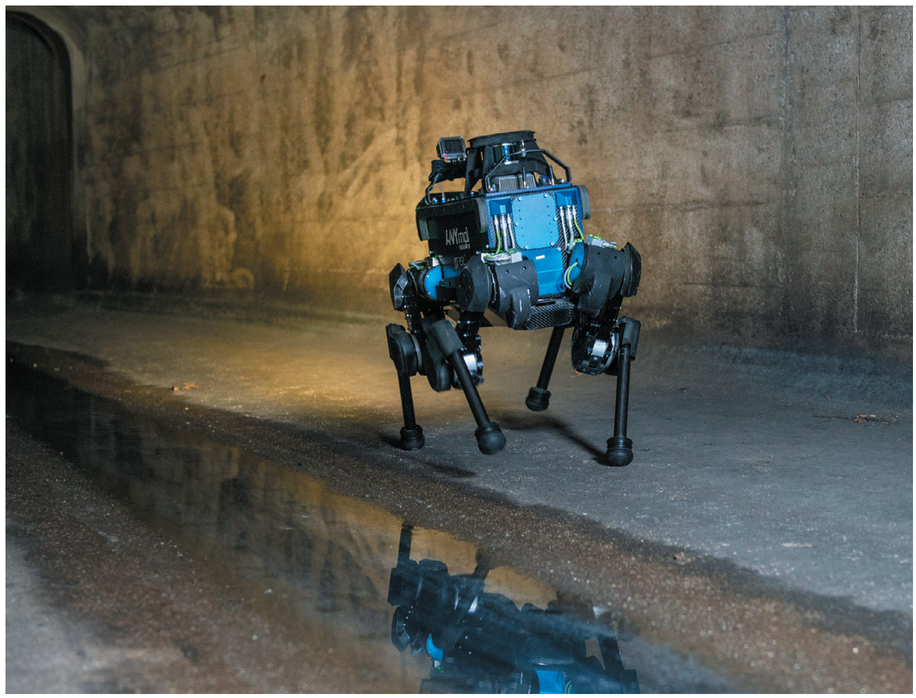

ETH ZURICH/DANIEL WINKLER

Figure 1 | The ANYmal robot. Hwangbo *et al.*[2] report that a data-driven approach to designing robotic software can improve the locomotion skills of robots. They demonstrate their method using the ANYmal robot — a medium-dog-sized quadrupedal system.

The hybrid simulator was faster and more accurate than a simulator that was based on analytical models. But more importantly, when a locomotion strategy was optimized in the hybrid simulator, and then transferred into the robot's body and tested in the physical world, it was as successful as it was in simulation. This long-overdue breakthrough signals the demise of the seemingly insurmountable simulation–reality gap.

The approach used by Hwangbo *et al.* hints at another major shift in the field of robotics. Hybrid models are the first step towards this change. The next step will be to retire analytical models altogether, in favour of machine-learning models that are trained using data collected from a robot's real-world environment. Such data-pure approaches — referred to as end-to-end training — are gaining momentum. Several innovative applications have already been reported, including articulated robotic arms[3], multi-fingered mechanical hands[4], drones[5] and even self-driving cars[6].

For now, roboticists are still learning to harness the power of faster computation, an abundance of sensor data and improvements in the quality of machine-learning algorithms. It is not yet clear whether it is time for universities to stop teaching classical control theory. However, I think that the writing is already on the wall: future roboticists will no longer tell robots how to walk. Instead, they will let robots learn on their own, using data that are collected from their own bodies.

Many challenges remain, of course, and chief among them is the challenge of scalability. So far, end-to-end training has been applied to physical robots that have only a small number of actuators. The fewer the actuators, the fewer the parameters that are needed to describe the robot's movements, and therefore the simpler the model is. The path to scalability will probably involve the use of more-hierarchical and modular machine-learning architectures. Further research needs to be done to learn whether end-to-end control can be scaled up to guide complex machines that have dozens of actuators, including humanoid robots[7], or large systems such as manufacturing plants or smart cities — urban areas that use digital technology to improve the lives of citizens.

Another challenge is less technical and more personal. For some researchers, the transition from using relatively simple mathematical models to applying 'black box' machine-learning systems — in which the internal workings are unknown — signals the unfortunate end of insight, and brings with it the feeling of loss of control. I am not one of those researchers. For me, there is something satisfying about seeing a robot, like a child, learn to walk on its own.

The insights offered by Hwangbo *et al.* could also be considered in the context of the mysteries of the mind. Consciousness has been one of the longest-standing puzzles of human nature[8]. In my experience, human-devised definitions of self-awareness are so vague that they are of little practical value for building robotic software. Perhaps the converse is true, however, and the study of robotic software can offer insights into age-old questions about the human mind.

One could conjecture that self-awareness and, by extension, consciousness are, at their core, an indication of our ability to think about ourselves in the abstract — to self-simulate. I would argue that the further ahead in time a person can look, and the more detailed the mental picture of their future activities is, the greater that person's capacity for self-awareness will be. Now, robots are capable of learning to self-simulate. This breakthrough is not merely a practical advance that will save some engineering effort, but also the beginning of an era of robot autonomy. ■

Hod Lipson *is in the Department of Mechanical Engineering and Data Science, Columbia University, New York, New York 10027, USA.*
e-mail: hod.lipson@columbia.edu

1. McNeill Alexander, R. *Principles of Animal Locomotion* (Princeton Univ. Press, 2003).
2. Hwangbo, J. *et al. Sci. Robot.* **4**, eaau5872 (2019).
3. Kwiatkowski, R. & Lipson, H. *Sci. Robot.* **4**, eaau9354 (2019).
4. OpenAI *et al.* Preprint at https://arxiv.org/abs/1808.00177 (2018).
5. Gandhi, D., Pinto, L. & Gupta, A. *IEEE Int. Conf. Intell. Robot. Syst.* **2017**, 3948–3955 (2017).
6. Bojarski, M. *et al.* Preprint at https://arxiv.org/abs/1604.07316 (2016).
7. Kuindersma, S. *et al. Auton. Robot.* **40**, 429–455 (2016).
8. Dehaene, S., Lau, H. & Kouider, S. *Science* **358**, 486–492 (2017).

This article was published online on 28 March 2019.

VIROLOGY

HIV remission achieved in the clinic again

A person infected with HIV who was treated for blood cancer with a stem-cell transplant has gone into viral remission, with no trace of the virus in their blood. A similar outcome in 2009 hadn't been replicated until now. **See Letter p.244**

TIMOTHY J. HENRICH

HIV infects immune cells, and the current standard treatment is long-term use of antiretroviral drugs. This keeps virus levels low in the bloodstream but doesn't eradicate HIV from cells in the body. In 2009, it was reported[1] that a person with HIV (commonly referred to as the Berlin patient) who was treated for cancer using a stem-cell transplant subsequently went into viral remission — the virus became undetectable in their body, even in the absence of antiretroviral therapy. No other cases of long-term HIV remission occurring in this way had been recorded since then. But now, on page 244, Gupta *et al.*[2] report a person who has achieved HIV remission for at least 18 months.

The case reported by Gupta and colleagues is similar in many ways to the one described a decade ago[1]. Both individuals had developed immune-cell cancer and received stem-cell transplants from donors (who were not infected with HIV) to re-establish their immune-cell populations (Fig. 1). Both donors had a mutation (termed Δ32) in both of their copies of a gene called *CCR5*. This gene encodes a receptor protein on immune cells that HIV can bind to during the process of infection. Having a Δ32 mutation in both cellular copies of the *CCR5* gene results in the absence of functional CCR5 protein on the cell surface, and immune cells lacking this protein can resist infection by HIV strains that depend on CCR5. Both patients were infected with HIV strains that exclusively use CCR5, along with the protein CD4, to aid cellular entry[3], which was probably a key factor in explaining why stem-cell treatment resulted in HIV remission.

In some HIV strains, the virus gains entry to cells by using a different protein in addition to CD4, typically CXCR4, rather than CCR5. These HIV strains are often associated with drug resistance, or might arise if antiretroviral treatment starts later than normal in the course of infection. It has been reported[4,5] that a person infected with HIV who received a transplant of donor stem cells that had the Δ32 mutation in both copies of *CCR5* experienced a rapid rise in HIV levels in their bloodstream when they stopped antiretroviral treatment, owing to a pool of pre-existing virus that could use CXCR4 for viral entry.

In the case reported by Gupta and colleagues, and in the 2009 report, the success of the

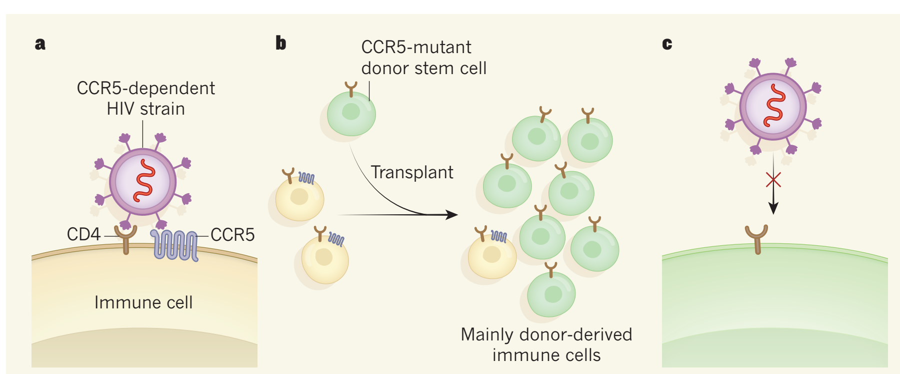

Figure 1 | HIV remission after cancer treatment. Gupta *et al.*[2] report the case of an HIV-infected person who underwent cancer treatment and then went into long-term HIV remission — the virus was undetectable in their bloodstream even in the absence of antiretroviral drug treatment — for at least18 months. This confirms an outcome[1] reported in 2009, which hadn't been repeated until now. **a**, The patient studied by Gupta and colleagues had a strain of HIV that depends on the CCR5 protein (along with the protein CD4) to attach to immune cells during infection. **b**, This person had treatment for blood cancer in which they received stem cells from a donor who had mutations that lead to the absence of functional CCR5 on the cell surface. The transplant resulted in the patient's immune cells being derived mainly from donor cells. **c**, The person's HIV went into remission, probably because the immune cells lacking CCR5 were resistant to infection by the HIV strain.