https://doi.org/10.1038/d41586-023-00710-0

News & views

Engineering

Swift progress for robots over complex terrain

Chen Li & Feifei Qian

A four-legged robot has learnt to run on sand at a faster pace than humans jog on solid ground. With low energy use and few failures, this rapid robot shows the value of combining data-driven learning with accurate, yet simple, models.

Anyone who has walked along a sandy beach knows how hard it is to move on sand. Like other granular materials, including mud and snow[1], sand yields and flows under the feet until they sink deep enough, and then it stops flowing and provides a stable foothold. In addition, sand doesn't spring back after impact, and the weight it can support before giving way depends on how wet and tightly packed it is, thus changing how much our foot sinks in and slips as we walk[1]. These complexities complicate the task of controlling a robot so that it can run on sand. However, writing in *Science Robotics*, Choi *et al.*[2] have succeeded in doing so, enabling a four-legged robot to be fast, robust and energetically efficient on sand.

Legged robots have, for several decades, been able to run on solid ground[3–5], and some robots that are small enough to fit in the palm of the hand have even done so on uniform sand in the laboratory[6]. Larger legged robots can walk slowly on natural granular materials[7,8], but researchers have struggled to control legged robots such that they match an animal's running pace on sand. Choi *et al.* managed this feat – achieving a top running speed of 3.03 metres per second – by integrating three approaches.

First, they used reinforcement learning[8] to train their robot to maximize its running speed and minimize how often it fails and the energy it expends. To do so, they first applied a technique called privileged learning, which is akin to training a teacher so that they can teach a student efficiently[9]. A simulated robot – the teacher – first trains itself to identify optimal control strategies by learning from a very large data set, which takes a long time. The student – the real robot – then benefits from what the teacher has already learnt, and can use partial, noisy data to quickly shift between control strategies. In the authors' case, the teacher learnt how to run under different sandy conditions in simulations, so that the student could adapt as it ran across real sand.

Second, to bridge the gap between simulation and reality, Choi *et al.* trained their robot teacher by simulating sand with highly variable physical properties and load-bearing abilities, similar to those found in nature (dry to wet, loosely to tightly packed). This is important because machine-vision systems, which are designed to see and interpret the world as eyes do, cannot reliably estimate the physical properties of a challenging terrain. For example, machine-vision systems might erroneously classify the top layer of wet sand as dry. Because dry sand flows more easily than wet sand, this error in judgement will affect the robot's performance. By exposing the robot to the variable physical properties of sand, Choi *et al.* improved the robot's ability to adapt to different sand conditions (Fig. 1).

Finally, to train their simulated robot, the authors selected and refined a model that describes the reaction forces exerted by sand on their robot's small feet as they strike it[10]. This model was not only accurate enough to capture the interaction physics, but also simple enough to make their simulations fast, both of which were necessary for the authors' success.

Aside from achieving their robot's remarkable performance and robustness, Choi and colleagues' work is notable for how well it integrates machine learning with models[11]. As more and more engineers aspire to ever-greater robot mobility in the real world using machine learning and simulations, these authors still appreciate the value of models developed through rigorous experimental

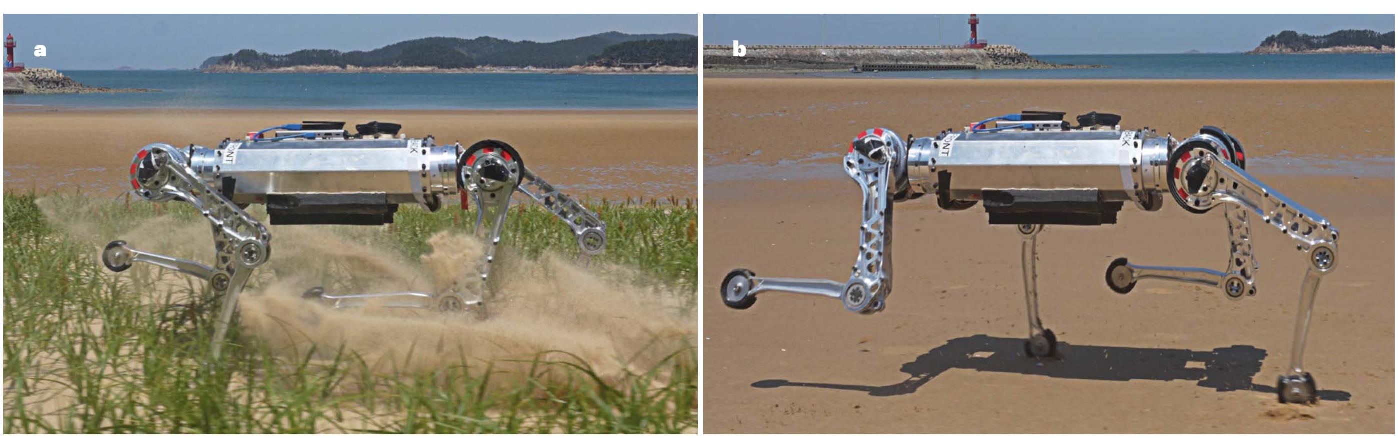

Figure 1 | A robot learns to run on sand. Choi *et al.*[2] used reinforcement learning to train a legged robot to run on level beach sand at a top speed of 3.03 metres per second. The robot's rapid, robust and efficient running was achieved by training it to adapt to variable terrain conditions, including dry and loosely packed sand (**a**) and wet and tightly packed sand (**b**). The training integrated data-driven learning and simulation with accurate, yet simple, models derived from rigorous experimental research. (Adapted from Fig. 1 of ref. 2.)

© 2023 Springer Nature Limited. All rights reserved.

research. They clearly went to great lengths to review the growing literature on the forces exerted by objects moving in sand; much of this research comes from laboratory experiments involving animal and robot locomotion[12].

This knowledge enabled Choi *et al.* to understand the nuanced contributions to forces on a small foot as it rapidly strikes sand[10]. One such contribution comes from the friction and pressure exerted by the weight of the sand particles. Another is a dynamic contribution due to particle inertia, which is similar to the aerodynamic or hydrodynamic drag felt in fluids. A third factor — transient, yet large — also comes from particle inertia, resulting from the sudden acceleration of the sand grains when they are first hit by the foot. The authors demonstrated that their refined model (which includes all these contributions) allowed the robot to achieve much higher performance and robustness than is possible with less accurate (and still often used) models.

Further research will be needed to improve legged robots so that they have animal-like mobility on terrain that is even more challenging than beach sand. As Choi and colleagues' robot moved on beach sand, only its feet sank — a scenario that is well described by the simple model that the authors used. However, achieving such mobility would be much more difficult if the robot's legs were to sink deeper into the sand[13,14] — for example, if it were much heavier itself[8] or if it were carrying a person or a large package. The robot would also face difficulties if it kept stepping into sand that was already disturbed[13], because the load-bearing abilities of the sand would continuously change. It would similarly struggle to move on steep dunes, which avalanche when disturbed[15]. Enabling robots to learn to deal with these extreme situations requires training them using models that capture the complex behaviours of sand, such as that described in refs 1 and 16, or models yet to be developed.

Choi and colleagues' demonstration is especially valuable and timely, given how pervasive data-driven learning approaches are becoming. These techniques have been successful at solving problems such as image classification, medical diagnosis, natural-language generation and game playing. But the authors' work is a reminder that good models are just as essential as data-driven learning for tackling problems such as robot locomotion in complex terrains. In many-particle systems such as sand, phenomena emerge that cannot easily be reconstructed from the fundamental laws of nature[17]. This calls for basic experimental research into the underlying principles and mechanisms, which machine learning might miss[11], and which simulations that are not vigorously validated by experiments might not fully capture[12].

Aerial and underwater vehicles that are autonomous, safe, fast and efficient have been engineered successfully because we have a fundamental understanding of aerodynamics and hydrodynamics. Choi *et al.* have set an excellent example by showing how the same level of success can be achieved for animal-like robots traversing natural terrains, building on the foundation of the emerging field of terradynamics[1]. By integrating fundamental research with data-driven learning and simulation approaches[11] in this way, we foresee exciting and rapid progress towards the goal of developing robots that can move across all terrains.

Chen Li is in the Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218, USA.
Feifei Qian is in the Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089, USA.
e-mails: chen.li@jhu.edu; feifeiqi@usc.edu

1. Li, C., Zhang, T. & Goldman, D. I. *Science* **339**, 1408–1412 (2013).
2. Choi, S. *et al. Sci. Robot.* **8**, eade2256 (2023).
3. Raibert, M. H. *Legged Robots That Balance* (MIT Press, 1986).
4. Saranli, U., Buehler, M. & Koditschek, D. E. *Int. J. Robot. Res.* **20**, 616–631 (2001).
5. Park, H.-W., Wensing, P. M. & Kim, S. *Int. J. Robot. Res.* **36**, 167–192 (2017).
6. Li, C. *et al. Proc. SPIE* **7679**, 76790Z (2010).
7. Raibert, M., Blankespoor, K., Nelson, G. & Playter, R. B. *IFAC Proc.* **41**, 10822–10825 (2008).
8. Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V. & Hutter, M. *Sci. Robot.* **5**, eabc5986 (2020).
9. Vapnik, V. & Vashist, A. *Neural Netw.* **22**, 544–557 (2009).
10. Aguilar, J. & Goldman, D. I. *Nature Phys.* **12**, 278–283 (2016).
11. Alber, M. *et al. npj Digit. Med.* **2**, 115 (2019).
12. Aguilar, J. *et al. Rep. Prog. Phys.* **79**, 110001 (2016).
13. Li, C., Umbanhowar, P. B., Komsuoglu, H., Koditschek, D. E. & Goldman, D. I. *Proc. Natl Acad. Sci. USA* **106**, 3029–3034 (2009).
14. Qian, F. *et al. Bioinspir. Biomim.* **10**, 056014 (2015).
15. Shrivastava, S. *et al. Sci. Robot.* **5**, eaba3499 (2020).
16. Agarwal, S., Karsai, A., Goldman, D. I. & Kamrin, K. *Sci. Adv.* **7**, eabe0631 (2021).
17. Anderson, P. W. *Science* **177**, 393–396 (1972).

The authors declare no competing interests.

© 2023 Springer Nature Limited. All rights reserved.