ARTIFICIAL INTELLIGENCE

Robot learning—Beyond imitation

Learning is a key topic for robotics. With the emergence of new learning methods and the maturity of existing approaches, extensive effort has been directed toward learning beyond simple kinematic planning, environment interaction, and elementary behavioral imitation from human demonstrators. In this issue, we have included a collection of papers covering different aspects and applications of robot learning, outlining current progress and opportunities, as well as the challenges ahead.

Training of sophisticated, dynamically balancing systems has been challenging, and the associated costs due to failures are a major limiting factor. In the paper by Hwangbo *et al.* (*1*), the authors used simulation data to train artificial neural networks and then transferred the networks to a legged (quadrupedal) robot for mastering dynamic motor skills. The end result is an agile, energy-efficient robot that can recover from falls even in complex configurations. Such sim-to-real transfer-based learning can be used in a diverse range of applications where real-life hardware-generated training datasets are scarce.

While humans seem to be able to pick or grasp objects of arbitrary shapes with ease, universal picking for robots remains difficult. In the paper by Mahler *et al.* (*2*), the authors present their Dex-Net 4.0, which can learn deep composite robot-grasping policies from large, physical-based synthetic training datasets—yet another example of how simulation-based learning transfer can be used to deal with real or novel objects. In (*3*), Ficuciello *et al.* investigated the problem of grasping novel objects with an anthropomorphic robotic system using synergistic control, demonstrating the practical use of combining learning from demonstration and reinforcement learning for complex tasks.

How to understand concepts and then generalize them to novel scenes is an important aspect of learning. Lázaro-Gredilla *et al.* (*4*) investigated how robots represent and infer high-level concepts by allowing so-called "zero-shot task transfer" to take place with inherently interpretable representations. They demonstrate that a robot could use their proposed method to interpret novel concepts (as schematic images) and then apply them to real objects on a flat surface.

With increasing popularity of soft robotics, how to use machine learning to derive intrinsically nonlinear behavior of the robot, as well as for sensing and perception, has attracted extensive interests. Thuruthel *et al.* (*5*) have used redundant sensors embedded in a soft actuator combined with recurrent neural networks to model an unknown soft continuum actuator in real time. The system is shown to be robust to sensor nonlinearities and drift and able to estimate the applied forces while interacting with external objects, thus enabling the development of force and deformation models for soft robotics.

Humans rely on different sensory data to see, feel, touch, and act. Here, Fazeli *et al.* (*6*) used a temporal Bayesian hierarchical model combined with tactile and visual sensing to develop a robot learning model to play Jenga, a game requiring complex manipulation skills. By using inferred beliefs, the robot adjusts its behavior and game strategy in a way that the authors believe to be similar to humans in learning, inference, and planning for complex physical interactions.

In this issue, we have also included four Focus articles discussing different aspects of robot learning from user intention detection and self-modeling to lessons that can be learned from decision neuroscience. In (*7*), Kim *et al.* present a learning-based intention detection method by using an egocentric (first-person-view) camera connected to a wearable hand robots. In (*8*), Kwiatkowski *et al.*

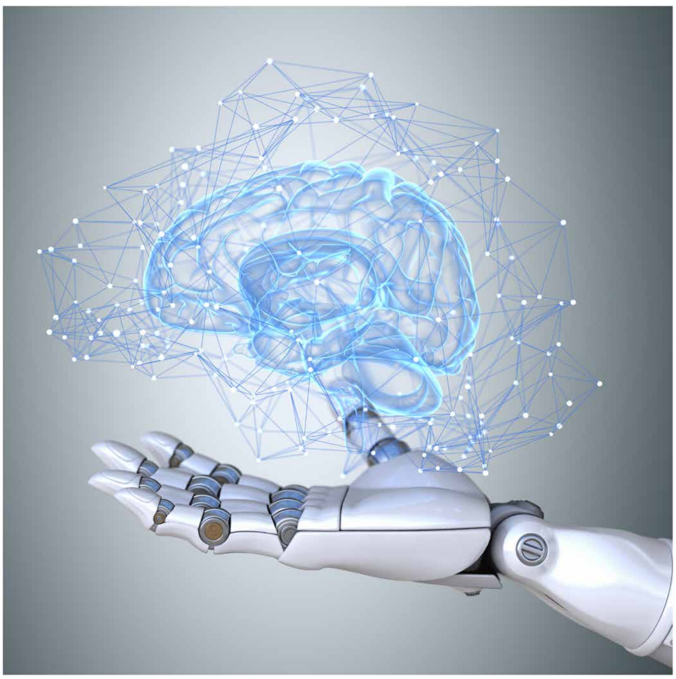

CREDIT: ILEXX/ISTOCK

Guang-Zhong Yang is the Editor of *Science Robotics* and Director and Co-founder of the Hamlyn Centre, Imperial College London, London, UK. Email: g.z.yang@imperial.ac.uk

outline how a robot can model itself without prior knowledge of its shape or physics (agnostic) and then perform specific tasks and even detect self-damage. The next two articles of this issue discuss, respectively, how robots reason about the meanings of human activities [purpose learning (*9*)] and how decision neuroscience can provide insights into the development of high-performance, memory-efficient reinforcement learning (*10*).

It is difficult, of course, for a single special issue to cover the broad range of research activities in robot learning. We hope that the examples provided here encourage much tighter collaboration between the machine learning and the robotics communities, driving toward ultimate goal of robot learning beyond imitation.

–Guang-Zhong Yang

REFERENCES

1. J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, M. Hutter, Learning agile and dynamic motor skills for legged robots. *Sci. Robot.* **4**, eaau5872 (2019).
2. J. Mahler, M. Matl, V. Satish, M. Danielczuk, W. DeRose, S. McKinley, K. Goldberg, Deep learning of ambidextrous robot grasping policies for universal picking. *Sci. Robot.* **4**, eaau4984 (2019).
3. F. Ficuciello, A. Migliozzi, G. Laudante, P. Falco, B. Siciliano, Vision-based grasp learning of an anthropomorphic hand-arm system in a synergy-based control framework. *Sci. Robot.* **4**, eaao4900 (2019).
4. M. Lázaro-Gredilla, D. Lin, J. Swaroop Guntupalli, D. George, Beyond imitation: Zero-shot task transfer on robots by learning concepts as cognitive programs. *Sci. Robot.* **4**, eaav3150 (2019).
5. T. Thuruthel, B. Shih, C. Laschi, M. Tolley, *Sci. Robot.* **4**, eaav1488 (2019).
6. N. Fazeli, M. Oller, J. Wu, Z. Wu, J. B. Tenenbaum, A. Rodriguez, *Sci. Robot.* **4**, eaav3123 (2019).
7. D. Kim, B. B. Kang, K. B. Kim, H. Choi, J. Ha, K.-J. Cho, S. Jo, Eyes are faster than hands: A soft wearable robot learns user intention from the egocentric view. *Sci. Robot.* **4**, eaav2949 (2019).
8. R. Kwiatkowski, H. Lipson, Task-agnostic self-modeling machines. *Sci. Robot.* **4**, eaau9354 (2019).
9. G. Cheng, K. Ramirez-Amaro, M. Beetz, Y. Kuniyoshi, Purposive learning: Robot reasoning about the meanings of human activities. *Sci. Robot.* **4**, eaav1530 (2019).
10. J. H. Lee, B. Seymour, J. Z. Leibo, S. J. An, S. W. Lee, Toward high-performance, memory-efficient, and fast reinforcement learning—Lessons from decision neuroscience. *Sci. Robot.* **4**, eaav2975 (2019).

10.1126/scirobotics.aaw3520

Citation: G.-Z. Yang, Robot learning—Beyond imitation. *Sci. Robot.* **4**, eaaw3520 (2019).

Science Robotics

Robot learning—Beyond imitation

Guang-Zhong Yang

Sci. Robot. **4** (26), eaaw3520. DOI: 10.1126/scirobotics.aaw3520

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.