TEACHING ROBOTS

ROBOTS LEARN TO LEARN, AND A BIT LIKE OUR KIDS DO IT

Spotlight Series

June 2022

FII INSTITUTE
Future Investment Initiative Institute
Impact on Humanity

TEACHING ROBOTS

ROBOTS LEARN TO LEARN – A LITTLE LIKE THE WAY OUR KIDS DO

Can machines learn to walk and act in the real world in a similar way to how children do it? Starting from scratch, watching, copying, trying, failing, retrying? Learning to move is not as easy as learning to think – but the robots are getting closer.

THE ISSUE AT STAKE

MYON WOULD BE A 6TH GRADER BY NOW. It was built in 2011 by Manfred Hild, a professor of neurorobotics in Berlin.[1] A humanoid-looking robot, Myon was specifically designed for – nothing. It had no special skills, but with 200 sensors, 50 motors and lots of limbs and joints, the machine was supposed to develop itself through observation of its environment.

"Myon has no goal," said its creator Hild. "But we do have a goal: to understand things." In this case, to understand learning. "Intelligence does not come overnight. Not in children, and not in robots either. How long does it take for children to stand, to walk, to talk?"[2]

So, instead of a program, the new robot was simply equipped with some rules to follow. For example, follow conspicuous signals. Or, once you've decided on something, stick with it for a while. And it was given a rather childlike design, with one overdimensioned eye in its head. The cuddle factor was supposed to lower the communication barrier and increase the patience of interacting humans. Look, it's a child; it's still learning.

Its moment of glory came in summer of 2015: Myon became an actor at the opera. At the Komische Oper Berlin, the Robaby played a leading role in the show "My Square Lady." A combination of video recordings and stage action allowed the audience to participate in the project and witness Myon's first successes, as well as its setbacks.[3]

Mainly the latter. The "follow the signal" rule, for example, caused Myon to turn its head to the loudspeakers, not to the singer. We know that a soprano's voice belongs to the person currently moving on stage, but the robot attributed it to the place the sound was coming from.

And Myon's performance hasn't improved much since then, as Benjamin Panreck, one of its co-creators, admits.[4] The robot needs intensive training to learn specific actions, and shows little sign of being able to transfer a once-learned lesson into a slightly altered situation. Today Myon is often used to teach students at professor Hild's neurorobotics lab how to train robots, but its own learning curve remains shallow at best.

ROBOTS CAN LEARN THROUGH INCENTIVES

But wait a minute, don't we live in the Age of Machine Learning? Yes, we do. It started 25 years ago, when Jürgen Schmidhuber and Sepp Hochreiter published their groundbreaking work on Long Short-Term Memory (LSTM; see box page 3).[5] They found a way for neuronal networks to "forget" certain outcomes and focus on the right ones. With sufficient training data, these recurrent neural networks can produce astonishing results for tasks like speech recognition→

COVER:

→ or machine translation. And they can do it almost autonomously, with relatively little human training effort involved.

Autonomously learning neural networks have become quite common. Autonomous robots, such as self-driving cars, are already a familiar concept. But autonomously learning robots remain only an aspiration. Robots that learn how to move and to act in the real world still rely heavily on human intervention.[6]

So how do robots learn? The main robot education principle is one they share with children and babies: learning through incentives. The robot first behaves in random ways and then evaluates how these behaviors have worked. That can be done via feedback from the instructor, who tells the robot whether its actions were effective or not. The robot chooses the behavior that offers it the highest reward, and then turns to the next iteration. It applies a number of random variations to the chosen behavior and determines by trial and error which of the new behaviors is now the most successful, and so on.

This method is called "reinforced learning,"[7] and as long as you stay in the pure world of data, it's not that different from the "forget gate" designed by Schmidhuber and Hochreiter. But robots by definition come in touch with a physical world beyond pure data, and then things become different – and more complex.

Take a robot that is learning to walk. In many cases the outcome is falling, usually a non-desired outcome. And then the →

LSTM - the backbone of robot learning

Over the last 25 years, Long Short-Term Memory technology (LSTM) has become a backbone of AI technology. Most of its use cases, such as speech recognition or machine translation, are part of the purely digital world, with input, output and process involving some kind of data. But the method can also be used to train robots – as a special layer between the input layer and the dropout layer.[8] One example is an OpenAI system called Dactyl: a human-like robot hand that learns to manipulate physical objects with unprecedented dexterity. And with a little help from an LSTM layer.[9]

PHOTO:

DEGREES OF FREEDOM

The Degree of Freedom (DOF) of a body is the number of independent motions among its pieces. A free 3-dimensional object has six degrees of freedom: up/down, left/right, forward/backward and three rotations around its axes.

Some joints, such as a human wrist, can freely rotate and have three degrees of freedom. Others, like the elbow joint, only allow one kind of motion. The more of these independent motions a robot can make, the more flexible it becomes – but also the more complex.

HUMAN	DOF
BODY	244
HAND	27
ARM	7
ROBOTS	**DOF**
KENGORO (UNIVERSITY OF TOKYO)	174
ASIMO (HONDA)	57
ATLAS (BOSTON DYNAMICS)	28
PEPPER (SOFTBANK)	20
ANYMAL (ANYBOTICS)	18

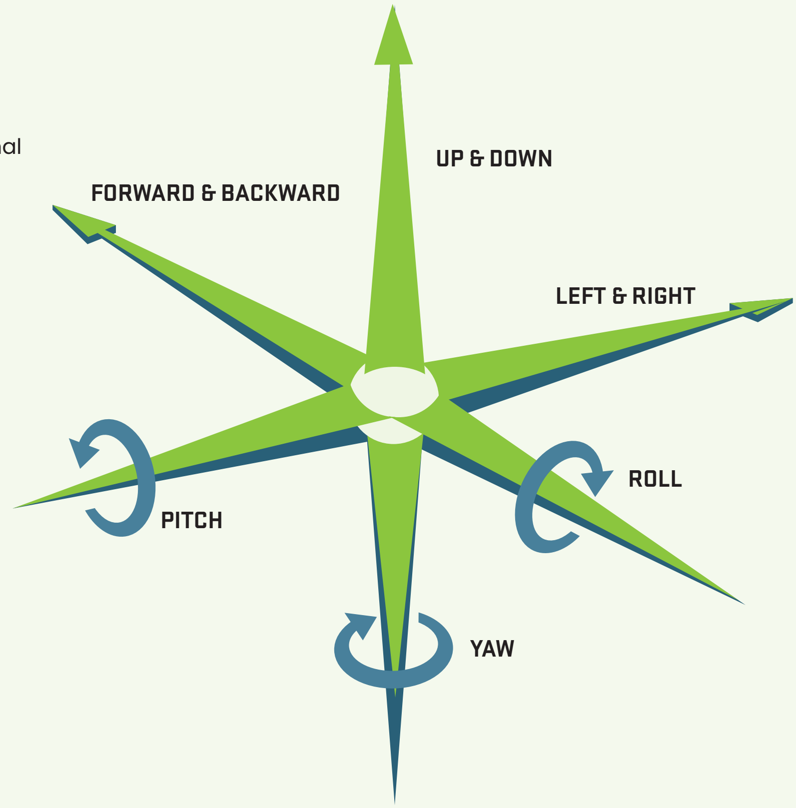

SOURCE: GUINNESS WORLD RECORDS[10]

robot should not simply forget the trial, but remember it so as not to do it again – like a child that touches a hot plate for the first time. So robo-learning through incentives can go both ways, reward and punishment.

And speaking of falling, that's costly. Every time the robot falls down or walks out of its training environment, it needs someone to pick it up and set it back on track. That's a lot of manpower, because robots need a lot of training. And it also poses a risk of damaging the robot. You may try to construct a robust robot, but to be able to walk, it also has to be flexible, with a lot of moving parts – joints and motors and sensors. And even if the damage risk is low for a single event, robots learning how to walk will produce a lot of such events during their training.

THE REAL CHALLENGE: THE REAL WORLD

When robo-learning is hard and risky, one of the ways to get ahead and improve the learning results is to reduce the complexity. The training ground can stay in the lab, or if you go out, you can at least stay away from hard-to-calculate surfaces, such as gravel or a doormat.

Sure, it makes the robot learn better – and reducing the variety of movements reduces the loss of time and material through falling and other failures. But without mastering real situations on the ground, the number of use cases for moving robots remains very limited – vacuum cleaning mainly, because Roombas don't need to walk. To get really productive out there, that's not enough. "A lot of places are built for humans, and we all have legs," says Jie Tan, Tech Lead Manager of the Robot Locomotion and Safety teams at Google Brain Robotics. "If robots cannot use legs, they cannot navigate the human world."[6] And that is the real challenge for robotics: the complexity of the real world.

The first ones seem to have made it. The robodogs from Boston Dynamics have been seen patrolling park lanes in Singapore and reminding citizens to maintain social distance.[11] They can also manage the slightly rougher terrain of the ancient Roman city of Pompeii.[12] But the task is still tricky. In Singapore, for example, the robodog did not walk autonomously (though it appeared to do so), but via remote control to avoid potential collisions with humans. And in the ruins of Pompeii, the robot only roams at night, when no visitors are around.

Right now, this turns out to be the silver bullet for robocompanies: stepping out into the real world, but only to places where no humans are around. That way, they can minimize irritation and risk, and gain room for learning (and failing). This approach has (up to now) only been tested at the DARPA Subterranean Challenge,[13] where autonomous robots had to solve a variety of tasks underground – in cellars, in tunnels, in →

PHOTO:

LEARNING THE HARD WAY

Robert Babuska, professor for cognitive robotics at the Technical University of Delft, Netherlands sees eight major challenges for learning robots in the real world.[7]

TOO MUCH COMPLEXITY

1. **Working in unstructured environments**
2. **Dealing with unforeseen circumstances**

TOO LITTLE REAL-WORLD EXPERIENCE

3. **Learning from few examples**
4. **Common sense (background knowledge of the world)**

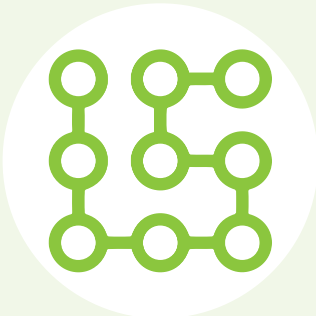

TOO LITTLE PATTERN RECOGNITION

5. **Generalizing learning ability to other circumstances**
6. **Using learning ability acquired in one domain in another domain (transversal learning)**

TOO MUCH MACHINE MINDSET

7. **Intelligent interaction with people**
8. **Creativity**

→ sewers. One part of the winning Team Cerberus were the ANYmal robots developed by researchers at ETH Zurich. These combine a visual perception of the environment with a sense of touch. Just like humans, they see and feel where they step, and they can relate this information to previous experience. So they can move quickly, robustly and safely over difficult terrain.[14]

The underground specialists are one of the best iterations on the way to learning to move. They are on track to replace humans in dangerous situations, such as on contaminated terrain, or in exhausting work, such as in agriculture (see box below). And they work mostly out of sight of humans. The Swiss robodogs of ANYmal, for example, are entering the market for the inspection of industrial sites, such as refineries for Indonesia's oil giant Petronas.[15]

NO NEED TO COMPETE WITH HUMANS

Though they are learning a lot, and continue to do so, it makes sense for robots not to compete directly with human workers. On the one hand, they would still probably lose the battle. On the other hand, their actions remain rather irritating for many humans. It also makes no sense for robots to become just like humans – humans have, for example, only two hands; robots can have as many hands as they want.

Knowledge transfer from one task to many others is an outstanding strength of →

Peeling bananas and picking strawberries

Fruits are as delicate to handle as they are delicious to eat. They come in different forms, sizes and stages of ripeness, and once you bruise them, you can't sell them. Challenge accepted: Using a method called "deep imitation learning,"[18] researchers at the University of Tokyo have taught a robot how to peel bananas. And there's even a fully autonomous strawberry harvester. Its 16 independently working robots are equipped with machine-learning vision systems that can determine when each fruit is ripe and ready for picking.[19]

→ the human brain - and a tough challenge for the neuronal networks of a digital brain. There has been promising research that one day may lead to "general purpose robots" – but it's still a long way from leaving the lab and succeeding in the real world.[16]

In the meantime, most robots succeed by simply not learning anything. They just do what they were designed for – with high speed and high precision, all day long and always the same. A welding robot in a car factory shouldn't try to learn painting, nor experiment to find the best welding spot; that's the engineers' job.

Or sometimes it's the artist's job: Swedish composer Fredrik Gran has programmed two industrial robot arms to play the cello.[17] He has not trained the moves, but has predefined every single move that they make; and this time, the outcome is not a car, it is music. The robo-cellists can play a composition – but not their own.

"Application-oriented systems can become perfect in what they were built and programmed for," admits Myon creator Hild. But, he adds, "this does not make them intelligent - and they cannot become intelligent in this way. For everyday intelligence to develop, you have to let a robot off the leash sometimes."

WHAT WE NEED TO LEARN

Or maybe we just have to accept that for quite some time we won't see the development of an everyday intelligence. For Hild, a simple example is basic household work. "If Myon should be able to help us at home, he would have to be extremely intelligent for that. Unless, of course, we're talking about cases like a washing machine, where everything runs according to a set program. But if a robot is to help in the household beyond those programmed actions, it needs a high level of intelligence."[2] And such a level of intelligence it nowhere in sight – not even close.

For Dutch roboticist Robert Babuska, this is mostly a problem of perception. The public expects too much and the experts are feeding the hype. "There's too much overselling," he says. "As roboticists we have to be more honest about what robots can and cannot do. Every time we speak to journalists or give a lecture, we have to clarify how complex it is to get a robot to operate well in the world."[7]

This, at least, is one task where Myon has always excelled: demonstrating how incredibly difficult it is to teach a robot.

←

WHAT CAN YOU DO?

➲ **See real-world robots as tools, not as toys:** Free-moving robots have reached the threshold of productivity. They should be used not for fun, but to replace human labor in dangerous or unpleasant environments.

➲ **Allow for mistakes:** Especially for learning robots, real-world complexity is hard to swallow. Assign them to tasks that offer room for failure and improvement.

➲ **Don't hold your breath for human-like robots:** The general purpose robot that can interact at eye level with humans is still far away. Technological development is moving toward specialized machines that excel in specific tasks – but do not compete with us.

PHOTO:

ABOUT FII INSTITUTE

THE FUTURE INVESTMENT INITIATIVE (FII) INSTITUTE is a new global nonprofit foundation with an investment arm and one agenda: Impact on Humanity.

Global, inclusive and committed to Environmental, Social and Governance (ESG) principles, we foster great minds from around the world and turn ideas into real-world solutions in five critical areas: Artificial Intelligence (AI) and Robotics, Education, Healthcare and Sustainability. We are in the right place at the right time: when decision-makers, investors and an engaged generation of youth come together in aspiration, energized and ready for change.

We harness that energy into three pillars: THINK, XCHANGE, ACT. Our THINK pillar empowers the world's brightest minds to identify technological solutions to the most pressing issues facing humanity. Our XCHANGE pillar builds inclusive platforms for international dialogue, knowledge-sharing and partnership. Our ACT pillar curates and invests directly in the technologies of the future to secure sustainable real-world solutions. Join us to own, co-create and actualize a brighter, more sustainable future for humanity. ←

FII INSTITUTE
Future Investment Initiative Institute

Contact

FII Institute:
THINK
think@fii-institute.org

References

1 Manfred Hild et al.: The Distributed Architecture for Large Neural Networks of the Humanoid Robot MYON. International Conference on Neural Computation Theory and Applications, 2011, [PDF]

2 Manfred Hild: Baby statt Sklave, in GDI Impuls, December 2015 (in German)

3 Kulturstiftung des Bundes: My Square Lady, 2015 www.kulturstiftung-des-bundes.de/en/programmes_projects/theatre_and_movement/detail/my_square_lady.html

4 Interview for this publication, May 2022

5 Sepp Hochreiter,Jürgen Schmidhuber: "Long short-term memory", 1997 doi.org/10.1162/neco.1997.9.8.1735

6 Karen Hao: "This robot taught itself to walk entirely on its own", Technology Review, March 2020 www.technologyreview.com/2020/03/02/905593/ai-robot-learns-to-walk-autonomously-reinforcement-learning/

7 Robert Babuska: Robots that learn like humans, TU Delft www.tudelft.nl/en/3me/research/check-out-our-science/robots-that-learn-like-humans

8 Raj Patel et al. Deep Learning-Based Robot Control using Recurrent Neural Networks and Adaptive Sliding Mode Control, Proceedings of CDSR'21, May 2021, [PDF]

9 OpenAI: Learning Dexterity, July 2018 https://openai.com/blog/learning-dexterity

10 Guinness World Records https://www.guinnessworldrecords.com/world-records/558450-most-degrees-of-freedom-for-a-humanoid-robot

11 GovTech Singapore, press release, May 2020 www.tech.gov.sg/media/media-releases/spot-robot-trial-for-safe-distancing-operations

12 Emmett Smith: A robot dog now guards the ancient streets of Pompeii, mashable.com, March 2022 https://mashable.com/video/pompeii-boston-dynamics-spot

13 DARPA Subterranean Challenge, September 2021, https://subtchallenge.com/index.html

14 Team Cerberus, https://www.subt-cerberus.org/

15 Damon Evans: Petronas tests explosion-proof robot dog offshore, Energy Voice, February 2021 www.energyvoice.com/oilandgas/300076/petronas-tests-explosion-proof-robot-dog-offshore/

16 OpenAi: Robots that learn, May 2017, https://openai.com/blog/robots-that-learn/

17 Fredrik Gran: Cello Suite No. 1, https://fredrikgran.com/works/cello-suite-no-1/?o=quintet

18 Emmett Smith: A delicate robot learned how to peel, mashable.com, April 2022 https://mashable.com/video/banana-peeling-robot

19 Harvest CROO, https://www.harvestcroorobotics.com/