

Robots That Learn!

In the world of science fiction, robots are intelligent and adaptive, but reality differs significantly. Robot programming is expensive manual labor, and the resulting programs are inflexible. A key step in making current robots more like their sci-fi counterparts requires endowing them with the capability to learn how to react appropriately and at the right time. **Jan Peters** is trying to teach exactly this skill to his machines. The computer scientist and mechanical and electrical engineer heads up a research group at the **Max Planck Institute of Biological Cybernetics** in Tübingen.

TEXT **TIM SCHRÖDER**

Jan Peters' laboratory is tightly crammed with equipment – hardly surprising, as the center of the room is occupied by a ping-pong table. The first thing one sees on entering is a heavy industrial robot – an articulated arm, as tall as a man and as thick as a sumo wrestler's thigh. A second robot arm is suspended from the ceiling next to the ping-pong table, and appears lighter and more dexterous.

Peters and his four students have pushed their desks over to the windows. Unlike the robots, neither the scientists nor their computers require much space. Peters is teaching his robots forehand and backhand, games of skill and how to grasp unfamiliar objects – all things that are still difficult to teach a machine, despite the fact that programmable robots have been around for half a century.

There is no question that modern robots can perform a lot of tasks. They weld auto body parts together with millimeter precision a thousand times a day without getting tired. Robot dogs play football, flick the ball into

the goal with their nose and roll yapping onto their backs when they score. Humanoid robots move forward on two legs, offer their hand and say “Hi.” And yet the robot world is still as rigid and inflexible as the stone circle at Stonehenge.

IMPROVED FLEXIBILITY AND APPROPRIATE REACTIONS


Factory robots repeat the same movement a thousand times a day, no more, no less. If a component rolls out of position, they can only report their switch into failure mode. To prevent them from injuring anyone, they are locked away behind bars and light barriers. Football dogs lose all sense of direction when the barriers are removed from the playing field. And electromechanical humanoids stumble when steered to walk over a fleecy velour carpet.

The fact that today's robots are still dumb is demonstrated by their inability to adapt to the uncertainty in human environments. If you affectionately slap your humanoid companion on the shoulder, you can expect

to be knocked to the floor, as the robot is incapable of processing an unexpected slap.

Greater flexibility and the ability to react appropriately and at the right time – this is what robots still have to learn, and precisely what Jan Peters is teaching his machines. Peters is a computer scientist as well as a mechanical and an electrical engineer. He heads a research group on robot learning in the Empirical Inference Department at the Max Planck Institute for Biological Cybernetics in Tübingen. Next door, his colleagues are busy developing theories and calculation specifications for statistical learning. Others employ these methods in brain-computer interfaces, computational photography and bioinformatics.

Jan Peters is the only one here whose laboratory contains massive robots mounted on the wall or suspended from the ceiling on concrete beams as thick as tree trunks. Whatever Peters and his students – computer scientists and engineers – come up with is sent straight to the machine and converted into movement. The Tübingen-based



No child's play for a robot: In order to catch the ball in the cup (blue), the arm first imitates the movement demonstrated by the scientists in Tübingen, and subsequently improves the behavior by trial and error.

» The Tübingen-based scientists are trying to teach robots how to learn. The basic idea is to teach the machines how to accomplish complex tasks by learning flexible motor primitives instead of employing manually programmed movements.

scientists are trying to teach robots how to learn. The basic idea is to teach the machines how to accomplish complex tasks by learning flexible motor primitives instead of employing manually programmed movements.

JOINING FORCES: ROBOTICS AND MACHINE LEARNING

Their approach is inspired by human motor learning: babies master the art of grasping things and, over time, learn to change their grip according to need. Grabbing quickly, taking hold of something firmly or gently – eventually, grown-ups master more than 60 types of grasps. But how to make a machine created from cable and steel as smart as an infant? The scientists in Tübingen have chosen to

bridge two disciplines that have co-existed for years: robotics and machine learning. In Peters' lab, the two approaches join forces.

Conventional robotics develops machines that are customized for precisely specified tasks, such as robots in an auto factory. The programmer determines a desired trajectory with high accuracy, as well as how much force is needed to accomplish the task and how it should move its joints. The result is a mechanical servant that strictly carries out the exact same orders. Machine learning is less concerned with heavy robot hardware and is located in the more virtual realm of "data clouds."


For example, computers learn to identify specific patterns. In mail distribution centers, character recognition

programs read addresses in fractions of a second. They have learned what the letter "A" looks like, and can assess very accurately whether an illegible letter is actually an "A" or some other letter. Peters combines the two disciplines to create a teachable robot. "Around 5,000 scientists worldwide are working on machine learning, and 6,000 or so more are engaged in research on robotics," says Peters. "But only six research groups are really combining both – and we're one of them."

A ROBOT LEARNS HOW TO CATCH

Why does machine learning not translate directly into robot learning? The answer becomes obvious when one watches a robot learning to play, for instance, ball-in-the-cup. In this game

Katharina Mülling guides the robot's arm, showing it how to hit the ball with the ping-pong paddle (left). Bowden cables on the robot's joints allow the arm to be made compliant, so that physical contact with it is quite safe (right).



Photos: Axel Griesch (2)



of skill, a string with a ball attached hangs from a small wooden cup in the robot's hand, and the ball is tossed into the cup by swinging it carefully to and fro.

Trying to teach a robot to play ball-in-the-cup through off-the-shelf machine learning methods would test the patience of a saint! Approaching the problem without domain insights from scratch requires testing thousands of variants, altering arm acceleration, joint rotation and direction of movement. "With this type of robot arm, there are so many permutations that it would take forever to achieve the desired results and catch the ball," says Jan Peters. Programming the task as in industrial robotics is no alternative due to the complexity and uncertainty of the swing movement. To start with,

the ball swings gently to and fro, before being spun around and flipped into the cup. The to and fro motion is very jerky and uneven. "It is precisely these complex movements that are so difficult to program using traditional methods," explains Peters.

LEARNING FROM IMITATION AND SUBSEQUENT TRIAL AND ERROR

He has therefore developed a robot learning system made up of several components. The learning and retrieval of movements is broken down into several easily digestible stages, using unusual methods such as imitation. First, you have to spell out to the robot what it is actually supposed to do. The quickest way to achieve this is to demonstrate the movement to the ma-

The robot is already doing well competing against a ball gun. However, before it can play with a human partner, it needs more practice. Katharina Mülling holds on to the emergency off-switch while Jan Peters and Jens Kober monitor its training.

» First you have to spell out to the robot what it is actually supposed to do. The quickest way to achieve this is to demonstrate the movement to the machine, in the same way that a tennis coach takes a pupil by the hand to practice his or her forehand.

chine, in the same way that a tennis coach takes a pupil by the hand to practice his or her forehand.



To play ball-in-the-cup, Peters' student Katharina Mülling guided the arm of WAM, the laboratory robot, swinging the ball into the cup. This demonstration gave WAM a rough idea

of the sequence of movements. It was then followed by a second step: self-improvement. Just like a human learning tennis, WAM had to perfect the movement in order to flip the ball into the cup all by itself. In the language of robot experts, this step is known as reinforcement learning. "Ultimately, the

robot's goal is to reduce the distance between the ball and the cup," explains Jan Peters. "If the ball ends up in the cup, the distance is minimal."

The experiment was a success: Once Katharina Mülling had shown WAM what to do, the robot needed only 45 attempts to learn a successful

A test of endurance for any waiter: Jan Peters and Duy Nguyen-Tuong nudge the robot's arm to teach it to compensate for their prodding, and to balance the bottle on the tray.



movement. After around 90 attempts, it now sinks the ball every time. According to Max Planck research scientist Peters, this is no mean feat. "We've had visitors here who were totally frustrated after trying unsuccessfully for much longer."

MAKING THE ROBOT COMPLIANT

The robot skill learning system in Tübingen goes even one step further. First, it stores the learned movements, the motor primitives. These are managed by a software module known as the supervisor. The idea is that, according to the situation, the supervisor gives the command to perform specific motor primitives. In most cases, the motor primitives must be adjusted to the new situation, or a series of learned motor primitives can be combined in order to react correctly in a matter of seconds – in the same way that a tennis player constantly has to vary his or her forehand moves: arm outstretched, volleying or diving for the ball. "Humans frequently experience situations in which they suddenly have to adapt a learned behavior," says Peters. "The first time they strap on a pair of inline skates, for example."

In the meantime, WAM no longer just plays ball-in-the-cup, but also ping-pong. It still practices using the reinforcement learning strategy. Jan Peters has screwed a ball gun to the ping-pong table, which delivers the balls more uniformly than a human could. WAM parries the shots easily. However, it cannot yet hold its own against a human opponent. After all, ping-pong is an incredibly fast sport, and also an extremely complex one. If a ball smashes onto the table or is



A matter of concentration: Manuel Gomez Rodriguez practices steering the robot arm using his brainwaves. Electrodes channel the commands straight from his brain.

whacked to the edge of it, WAM has to move its arm incredibly fast and position it correctly within a fraction of a second in order to hit it.

Unlike conventional robots, WAM robots are fitted with flexible Bowden cables, like those used in automobile brakes, instead of rigid electromechanical gears. These enable a WAM to gently follow trajectories and even give way when it collides with its human training partner – a key criterion in interaction between technology and humans. "If a Bowden cable breaks, we can easily spend a few hours repairing

it and fiddling around with screws. But such a compliant robot has huge advantages," says Peters.

MOTOR PRIMITIVES PROVIDE REAL TIME FLEXIBILITY

The scientists and WAM are currently learning what forces are needed and how fast the joints need to be moved for an action-packed game of ping-pong, how quickly the Bowden cables react and how to steer the arm correctly. They are still working on it. One thing is certain – the motor skills learn-



Mathematical instructions: Katharina Mülling and Jan Peters derive equations in order to teach the robot new tricks, while Jens Kober and Oliver Kroemer discuss the movements the robot needs to master in order to water flowers.

ing system will break a ping-pong game down into a series of motor primitives. This part seems to be the preferred solution in order to obtain real time flexibility. Any other elaborate calculations would take too long. A clean backhand also requires the robot to hit the ball quickly. WAM uses several cameras to track the ball's position. Fast-acting image processing software is therefore a must.

A REWARD FOR A "GOOD ROBOT"

Six working groups worldwide make for a manageable scientific community. Jan Peters and his students know the other research scientists in the US and Japan well. Peters worked for a long time at the University of Southern California in Los Angeles, where he earned his Ph.D. He uses the prin-

ciple of rewarding when referring to reinforcement learning, as used by Pavlov in his dog experiments. A successful attempt is rewarded with food. "The robot learns how to become a 'good robot,'" Peters explains.

Depending on the robot's goal, it can then adjust its motor primitives. Accordingly, the stored sequence of movements varies based on the goal function. Peters' student Jens Kober initially experimented with this virtually, teaching his computer to hit specific segments on an imaginary board with a dart. A few weeks ago he was in Japan and transmitted his commands to the laboratory robot. "Admittedly, it sometimes missed, because the mechanics of the robot hand didn't always cooperate. However, some darts landed exactly in the right corner," says Kober.

A supervisor, a motor primitive archive and goal functions are just some components of the robot learning system. Several other factors come into play before the command to move is given.

However, it is clear that all components interact to achieve the movement, and this part is achieved by distributing the task in manageable chunks to these different modules.

FLEXIBILITY IS NOT A REQUIREMENT

Service robots already exist that can be taught to move by means of instruction. However, they generally act only as a third hand, lifting and holding heavy sheet metal or engine parts. Ultra-fast flexibility is hardly a requirement.

“If service robots are actually going to help us with housework or nursing care in the future, they have to be able to do it properly and adjust quickly to changes in environmental conditions,” explains Peters. There appears to be a big demand for these service robots. Globally, five million iRobot vacuum cleaners are already whirring around people’s homes. But here again, flexibility is not a requirement. Now and then, cleaning robots will buzz through the living room for a half hour before gliding back to their charging stations. WAM would probably just give them a tired smile. ◀

www.robot-learning.de

GLOSSARY

Empirical inference

The Biological Cybernetics Department founded in 2001 at the Max Planck Institute of the same name researches the physical laws underlying empirical data. Scientists develop algorithms and apply them to a wide range of problems, such as computer vision, brain-computer interfaces, bioinformatics ... and robot learning.

Humanoid robot

The design of this type of robot mimics the human body. It normally has a head and two arms and legs. Movement sequences and joint positions are also based on those of humans.

Machine learning

An artificial system, in this case a robot, does not memorize all examples, but rather gradually acquires the ability to generalize. The system recognizes physical laws in the learning data and is eventually also able to evaluate unknown data or react to new situations.

Pavlov’s dog

Russian Nobel Prize winner Ivan Petrovich Pavlov (1849 to 1936) conducted the first empirical experiments to show the effect of classic conditioning: he observed the connection between salivation and digestion. The owner’s footsteps alone sufficed to trigger salivation in kenneled dogs, even when no food was being offered.