Aspiration beyond reality: It will be a while before robots can dance as smoothly as humans and also improvise movements like in this AI-generated photo. Giving them some kind of body awareness is a step in that direction.

IMAGE: AI IMAGE MIDJOURNEY | CREATED BY GESINE BORN | BILDERINSTITUT

ROBOTS DISCOVER THE WORLD

TEXT: ROLAND WENGENMAYR

Robots can already assist humans with some everyday tasks. But they are out of their depth when faced with unfamiliar environments or even small deviations in the tasks they are trained to perform. To help them learn to adapt more quickly to new circumstances, Michael Mühlebach and Jörg Stückler's research groups at the Max Planck Institute for Intelligent Systems in Tübingen are developing new training methods for the machines. Their robots even have to prove themselves by engaging in table tennis or body flight.

Intelligent robots were already a technological myth before there were machines that even remotely deserved the name. But what can robots really do today? How far are they from science fiction icons like the comically humanized C-3PO from Star Wars? A search on YouTube quickly leads to a video by the US robotics company Boston Dynamics. It shows the humanoid robot Atlas, who dazzles with its somersaults, runs and hops over a challenging training course together with a twin sibling, or supports a human on a scaffold. But as impressively light-footed and almost uncannily human-like as these robots' movements are, they are performing these movements in a familiar environment for which they have been trained. It is not publicly known what Atlas and its ilk would actually be capable of if they had to orient themselves in a completely new environment and act independently – this is a company secret.

Joachim Hertzberg, Scientific Director at the German Research Center for Artificial Intelligence and professor at the University of Osnabrück, is impressed by the complex ways in which Boston Dynamics' robots can move. But he also immediately mentions a major caveat: if you ordered one of today's robots to carry out a task on its own and follow a plan in an unfamiliar environment, even if it was only to fetch a coffee, the result would be considerably less spectacular. "The field is called Artificial Intelligence, but it comes down to intelligence that we ourselves regard as completely unintelligent," Hertzberg says, "that is, the ability to reasonably navigate an environment, to do tasks for the first time without having practiced them beforehand, to act in accordance with the situation and goal."

Flexible algorithms

Machines that are constantly learning – and learning as quickly as possible – are one step on the path towards robots that retain their orientation even in unfamiliar environments and when undertaking new tasks. Two teams at the Max Planck Institute for Intelligent Systems in Tübingen are working on such systems. Unlike companies or application-oriented research facilities, the researchers reduce the complexity of the tasks their machines have to master in order to first teach them elementary aspects of orientation. Michael Mühlebach's group is looking at how robots can train a kind of body awareness, to put it in human terms, through prior knowledge of their own physical properties. Without this, they will not be able to move in an unfamiliar environment and carry out commands precisely. This is also what is happening in Jörg Stückler's group, which is working on teaching robots to see. This seeing involves robots learning to recognize both stationary and moving objects in any environment. In this case, though, recognition is limited to purely physical properties of the objects, such as their size, shape and color, which alone is a formidable challenge.

A permanently learning algorithm is used, for example, in the table tennis robot Pamy, whose hardware was largely developed by Dieter Büchler's group at the Max Planck Institute in Tübingen. Pamy must react flexibly to changes in the game. For now, the one-armed robot is training with a ball machine to learn how to correctly estimate the future trajectory of a ball. The experiment takes place in a laboratory under Michael Mühlebach's leadership. There, his doctoral student Hao Ma welcomes us to a background of noise that sounds as if we had landed in training camp of an air pump team.

Hao Ma has to grin at the sight of the perplexed-looking guest and points to a cordoned off area. There, puffing loudly, a single robotic arm performs wild dry runs

PHOTO: WOLFRAM SCHEIBLE FOR MPG

SUMMARY

Researchers at Max Planck are creating techniques that will help machines recognize new objects and orient themselves more quickly in unknown and dynamic environments.

Simple physical models that give robots a head start on understanding their own movements, their surroundings, and the objects they interact with are one way they speed up the learning process.

For example, the robots can learn to play table tennis, float in the air, and deduce object properties from image data.

without a ball on a kind of platform with a ping-pong paddle. Two aluminum tubes, connected with plastic joints, form the upper and lower arm; the hand consists of a joint with a firmly screwed-on bat. Air hoses lead to the joints and are connected to a battery of pneumatic cylinders at the bottom of the platform. In the transparent cylinders, pistons can be seen pounding up and down. They push air into or suck it out of the pneumatically driven joints and thus move the arm.

Muscles driven by air pressure of this kind allow a very light-weight design without electric motors at the joints, which is why a robotic arm can perform fast movements. However, this design has one drawback: the arm visibly rebounds. To accurately return a table tennis ball, the controller needs to be intimately familiar with this elastic behavior. With the help of cameras, angle sensors, and pressure sensors that track its movements in real time, Pamy is currently learning the necessary body awareness by playing air ping-pong. This learning process is where a key research approach of Mühlebach's team comes into play. To save the arm's controller from having to painstakingly start from scratch when learning its properties, the team has already programmed Pamy with a straightforward physical model. It represents the arm with ideal stiff bars and idealized joints. "What is difficult to describe, however, is the behavior of the 'muscles' from the plastic containers that fill with air," Mühlebach explains: "I use machine learning to do this." The algorithm used for this purpose only uses camera recordings to learn the rebounding of the pneumatic movements. This saves a lot of computing time. Without the prior knowledge provided by the physical model, Pamy would need 16 hours to learn some kind body awareness, Muehlebach says. "With the model, we can get it done in about an hour." Accelerating robot learning through prior knowledge of physics is a key strategy in Mühlebach's and Stückler's research. Ma leads us to a ping pong table in the lab. There, the predecessor model of the robot arm, which is still busy playing air ping-pong, can now show what it can do. A rotating ball machine shoots a ping pong ball across the board, and it bounces once in Pamy's court before the machine knocks it cleanly back. With each new ball, it does this with impressive reliability. Ma points to four cameras mounted above the table. They follow the path of the bright orange balls. An algorithm has now learned to predict the future trajectory of a ball from the previous trajectory so accurately that the robot arm reacts like a skilled human player and hits it correctly.

One arm plays table tennis: Hao Ma puts the robot arm into position for a training session – the machine is tasked with developing a feeling for its own movements.

"In a new version of the ball prediction, we've also included how the ball is shot," says Jörg Stückler. This would be much more difficult with a human opponent, he says, but experience with the ball machine shows how it could work in principle. Pamy can also draw on prior knowledge about the ball machine. Jan Achterhold, a doctoral researcher in Jörg Stückler's team, taught this to the robot. The corresponding model even accounts for the fact that this machine has the ability to spin the ball. Once the ball has touched down in ro-

PHOTO: WOLFRAM SCHEIBLE FOR MPG

Know your opponent: The table tennis robot on whose control Jörg Stückler, Hao Ma and Michael Mühlebach (from left) are working plays against a ball machine. A model of this machine helps it to hit the fired balls.

bot arm's court, this results in it being deflected sideways. Pamy has to react to this immediately, which is a considerable challenge for the robot.

Body flight endurance test

Achterhold and Stückler used a gray-box model of the ball machine. Stückler explains that this is the halfway house between a black-box and a white-box model. The black-box model refers to machine learning without any prior knowledge, i.e., laborious trial and error. A white-box model would be the opposite: an alterable programmed physical model that is not capable of learning. In a simple, mechanically ideal world, this would work too, because the process, including the trajectory, could be calculated exactly. But with a real ball machine, effects always occur that the inflexible white-box model cannot deal with. Achterhold's team therefore employs a physically pre-trained machine learning system. To create this, the researchers first designed a physical model and combined it with a sophisticated learning algorithm that enables the system to learn the real properties of the ball machine. The team therefore used the advantages of both the black-box and white-box approaches. "That's why the approach is called gray-box," Stückler explains. Speaking with the robot researchers has made it increasingly clear what difficulties our seemingly routine, unconscious human behaviors pose for robotics. But this does not deter Michael Mühlebach from really wanting to know about it. "I'm fascinated by body flight, acrobatics, spins and tricks," he says, laughing. "And that's when I thought: wouldn't it be awesome to do this with robots!" In body flight, which is also called indoor skydiving, people float in a strong air current produced by a vertical upward wind tunnel. As in skydiving, they must learn how to specifically control their flight behavior on the cushion of air by changing their body posture and thus their aerodynamics.

In Tübingen, a lightweight flying robot, barely larger than a hand, will learn this above a mini wind tunnel. Doctoral researcher Ghadeer Elmkaiel is currently experimenting in his lab with his self-developed wind tunnel, which is designed to achieve the most uniform airflow possible using six propellers arranged in a circle. Above the opening of the wind tunnel is a holding device for the little flying robot. During training, it detaches from this device and tries to hover without connecting to a computer. In the process, it is supposed to gradually learn predefined flight maneuvers. It is not there yet, but again, prior knowledge of a simple physical model should speed up the learning process of the flying robot.

„It comes down to intelligence that we ourselves regard as completely unintelligent."

JOACHIM HERTZBERG

The knowledge Mühlebach's team gained from the acid test of the robots' orientation skills may also be used in unrelated fields, such as intelligent power grids. These are designed to match electricity production and distribution to demand as closely as possible. With the expansion of decentralized wind and solar power plants, whose electricity production also depends on the weather, this is becoming increasingly important and challenging. There are elements in such a network that can be accurately predicted by physical models as well as those whose behavior can only be predicted through experience. Large power plants that can have their electricity production physically modeled fall into the first category. The electricity market and the behavior of end consumers, on the other hand, can only be predicted through experience, for example, over seasons. "There's huge potential for machine learning there," Mühlebach says. But back to the orientation. Jörg Stückler's group is working, for example, on advancing a technology that combines camera data with the data from acceleration sensors, such as those built into smartphones. Stückler explains the acceleration sensors give a robot sense of balance, so to speak. By combining this with camera data, the robot should develop a knowledge of how its real body responds to a command. For example, if it is told to drive off, it needs a sense that it must first accelerate its mass to the specified speed. The robot develops this body awareness much faster if it has been programmed with a simple physical model of itself.

If a robot is commanded to handle objects, it not only needs a good feeling for its own movements, but also an idea of the objects and their properties. Granting the ability to recognize these through observations alone is the goal of doctoral researcher Michael Strecke. Since the camera data is noisy, i.e., blurred, it is not easy to read the shape or size of an object from it. However, if the robot observes how a repeatedly thrown ball dots against a wall, bounces back, and falls to the ground, it still learns something about the properties of the ball. It gradually understands how big the ball is, and that it therefore bounces back in a certain way. In this way, it learns to estimate how such an object is likely to behave through visualization alone.

In principle, therefore, it is possible to infer the properties of one object from the mere observation of its mechanical contact with another. This is how toddlers learn when they throw objects around and observe them. For computers with vision, this contact method has only worked so far for rigid objects, and also only for those with very simple geometry. Strecke and Stückler have now succeeded in advancing the machine learning of more complex shapes with a new optimization method. They illustrate this with a somewhat absurd example. A machine observes one object falling on another and initially mistakes it for a cow. Over the course of several collisions between the two objects, the machine gradually perceives the cow transforming into a kind of "rubber duck swim ring," and it falls onto a stick with its hole in the middle, like in a throwing ring game. This seemingly far-fetched scenario corresponds to a situation in which people also first have to completely reorient themselves. Robots are still in their infancy at this point, roughly at the same stage as a toddler, for whom every object in its environment is brand new. Using their new training methods, Jörg Stückler and Michael Mühlebach want to help the machines to orient themselves more quickly in unknown situations. But the road to a C-3PO complaining to his companion R2-D2 that they got into another mess, as in a Star Wars adventure, may still be quite long.

www.mpg.de/podcasts/orientierung (in German)