Robots
Go to School

As domestic help, healthcare assistants or emergency response units: robots are suitable for these jobs only if they are capable of learning and acting independently, at least to a certain extent. **Stefan Schaal** and the members of his Autonomous Motion Department at the **Max Planck Institute for Intelligent Systems** in Tübingen are teaching machines to become flexible and autonomous.

TEXT **PETER HERGERSBERG**

T ruth be told, Apollo doesn't exactly look like a Greek god. With that trusting gaze coming from his big, round eyes, he's more reminiscent of Shaun the Sheep than of the immortal being who struck fear into the hearts of his foes. At best, Apollo, from the laboratory of the Max Planck Institute for Intelligent Systems, could be compared to a demigod, considering that his upper body rests on a massive column rather than on a chiseled abdomen and legs. And the feats he accomplishes are, well, really the simplest of the divine exploits: if all goes well, he can securely grasp different objects, balance a rod on his hand or even mount a wheel on an axle.

That may not sound very impressive for a god, but the Apollo stationed in the laboratory of the Max Planck Institute for Intelligent Systems in Tübingen is a robot. And for a robot, he is actually capable of performing a surprisingly wide range of tasks. Perhaps most importantly of all, he learns a lot – and he does so in a way that might one day enable him, or rather his two-legged

The Director and his student: Apollo is one of the robots that Stefan Schaal and his team are teaching to move autonomously.

Photo: Wolfram Scheible

descendants, to autonomously move around in unfamiliar surroundings and independently solve difficult tasks.

Robots need to be able to learn new things if, for instance, they are to serve as emergency response units at a damaged industrial complex or rescue injured persons following an accident and are suddenly confronted with an unexpected obstacle. Also as domestic help or even as healthcare assistants, robots would need to be capable of continuously adapting to new situations and unforeseen events.

Stefan Schaal, Director at the Max Planck Institute in Tübingen, and his team have made it their mission to help these mechanically engineered beings achieve that degree of autonomy. To provide an idea of how much work still needs to be done in this field, Stefan Schaal shows some video footage taken at the Darpa Robotics Challenge. In this competition, robots must drive a golf cart and walk across sand to reach a building where they must close a valve. What we mostly see here, however, is rescue machines, accompanied by heroic-sounding music in the background, stumbling and falling down in various ways: as they get out of the golf cart, as they stagger across the sand, or as they attempt to push open the door and instead end up losing their balance and falling over.

The reason most modern-day robots falter in unknown terrain is because they have a one-track mind, in the truest sense of the word. An industrial robot that is programmed to screw doors onto a car chassis will execute that particular task perfectly – but only that one task, and only for a particular car model, at least as long as its software is not reprogrammed to carry out a new command.

And even the learning robots that already exist can only be taught one task each, for example by having a human trainer take and guide the robot's arm to help it carry out the desired motion, similar to what physiotherapists do with stroke patients in the early stages of rehabilitation. When the machine is then able to emulate a movement on its own, it receives positive or
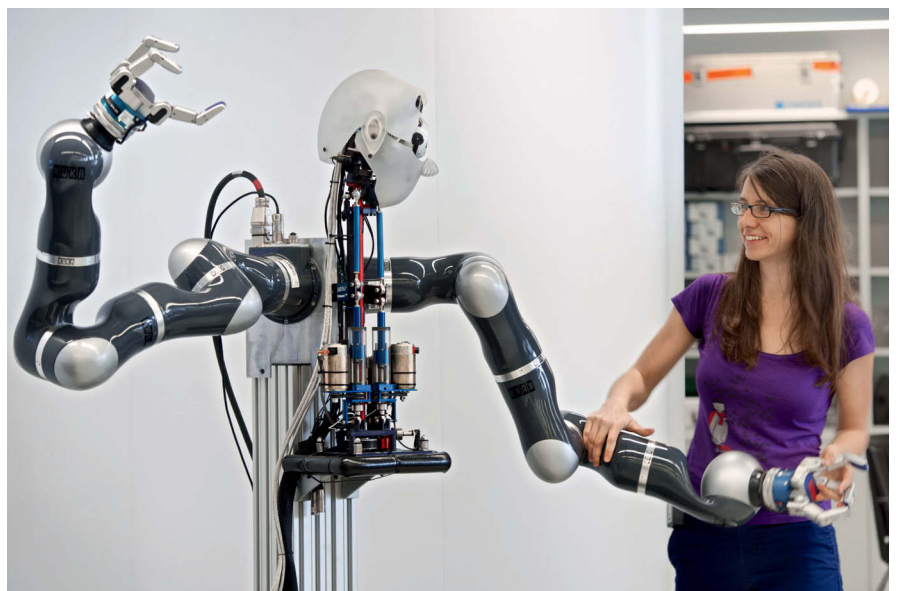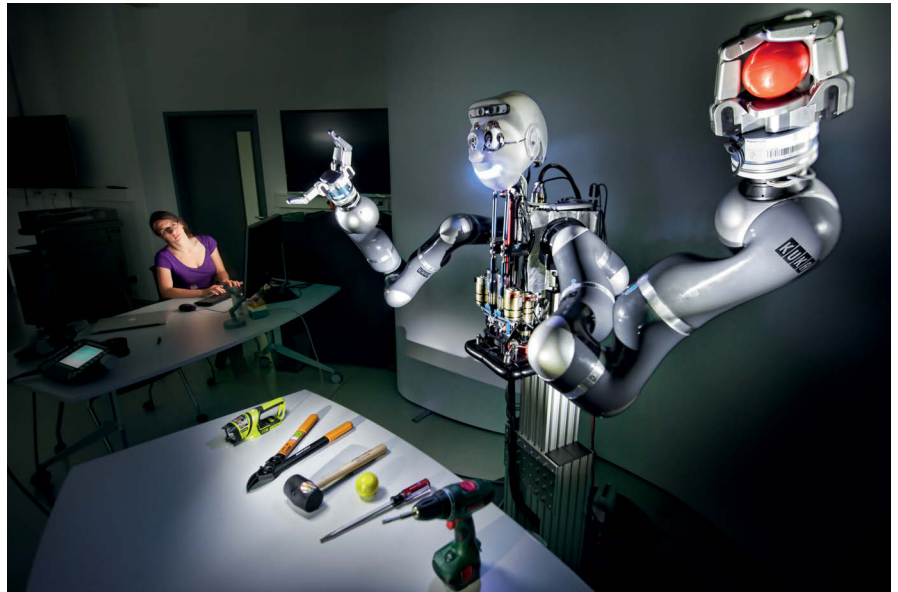
Coming to grips with it all: Jeannette Bohg teaches Apollo to plan the right way to grasp different objects based on visual input.

negative feedback until the robot can perform the task correctly, which it then does. But watch out if something unexpected interrupts it. "Today's robots are not robust," says Stefan Schaal. "They have a really hard time compensating for disruptions." If a robot learned to grab a hammer by the handle, for example, it would already consider it a disruption if it were handed the hammer head first.

"We want to achieve robustness by using various methods to introduce machine learning into the field of robotics," says Stefan Schaal, whose department focuses on perception-action-learning loops. When a machine – meaning a computer, which makes up the brain of every robot – learns something, a software program is trained to perform a particular task by feeding it large amounts of data. By using nu-

> **»** It would easily take a robot an entire lifetime to gain enough experience in order
> for it to become fully independent of human commands or intervention.

merous photos of people taken from different angles and in a variety of settings, for example, image recognition programs can be taught to reliably recognize faces – even when the latter are partially obscured or visible only in semi-profile.

This is the principle the researchers in Tübingen have been applying at their school for robots, which they have been running for almost three years now and which Apollo also attends. Yet the school could be considered more of an experimental educational institution because, unlike regular teachers, the Tübingen-based scientists don't teach their students existing knowledge, but rather start by determining what and how robots learn best.

One of the researchers involved in machine education is Jeannette Bohg. She trains the machines' visual perception in such a way that their visual sense provides them with the information they need for planning actions in a sensible manner. One of the goals, for example, is to teach the robots to analyze an unknown setting and then quickly and reliably find objects they need to solve a particular task.

When searching for a laptop, for example, software programs will use a bottom-up approach to look for conspicuous pixel clusters, or they will analyze all objects located in a given setting. However, computing all that information takes so much time that a robot can hardly complete the task within a reasonable period.

That is why Jeannette Bohg models her teaching on the top-down search strategy that humans use: "We know exactly where to look for something and where not to look for it because we have the necessary background information," the researcher explains. "If we were looking for a laptop, for example, we would expect to find it on a table, but not on a wall." A wall is where we might expect to find a clock, which, however, a person could also be wearing around his or her wrist as a watch. Looking at a scene and narrowing down the number of locations that would be worth searching is helpful for a robot, not least because it can then approach those locations and examine them more closely – just as we humans often do.

In order to teach her mechanical students these human search techniques, Jeannette Bohg researches the best way to model this human strategy with software. She then trains the software using the tracked eye movements of 15 participants who were asked to examine 400 pictures and look for a clock or a laptop, for example. This data allows the robot to acquire the experience that teaches a human where a particular object is most likely to occur.

ROBOTS WORLDWIDE COULD SHARE THEIR KNOWLEDGE

"Following the training sessions, our search algorithm is already quite good at locating clocks and laptops," says Jeannette Bohg. However, this technology is not quite as reliable at finding individual objects as methods that analyze the whole picture. "But bear in mind that, when using 400 pictures, the data set for the training is still rather limited," says Bohg.

Collecting sufficient data and drawing the right conclusion from it in order to be ready for all the contingencies of an autonomous existence is a general problem that machines face: Gaining enough useful experience to allow them to be independent of commands or interventions would easily take up a robot's entire lifetime, which is just as finite as the existence of a computer, a car or a human being. A single electronic brain would hardly be capable of processing such an enormous volume of data. "We might be able to solve this problem using cloud robotics," says Stefan Schaal. Similar to the way in which countless computers are already linked to solve large tasks, robots the world over could unite to altruistically share their knowledge – provided, of course, that their programs are compatible.

For now, each robot is left to its own devices to gather and process all the knowledge it needs to act in a halfway independent manner – for instance to correctly plan which grip to use when they see a particular object. This is another of Jeannette Bohg's areas of research.

In the past, robotics researchers programmed robots in such a way that the robot would first compute which points of an object its fingers needed to touch in order to grasp it securely. "Researchers proceeded on the assumption that the robot had internalized a detailed geometric model of both itself and the object, allowing it to compute and precisely reach the right points on the object in order to grab hold of it," explains Jeannette Bohg. The robot then used these models to plan how to grasp the object without it falling to the floor.

"But it turned out that these assumptions aren't realistic," says Jeannette Bohg. Not only because a robot's software doesn't include a model for each and every thing it could possibly grasp, but also because its controls weren't yet precise enough to reach the computed points on the object, especially since the data coming from the sensors it uses to control its movements is often incomplete and noisy. As a result, machines would often clumsily and unsuccessfully try to get hold of an object. Jeannette Bohg wants to change that, and once again her work takes its cues from humans, who are capable of reliably grasping even objects they have never seen before. >
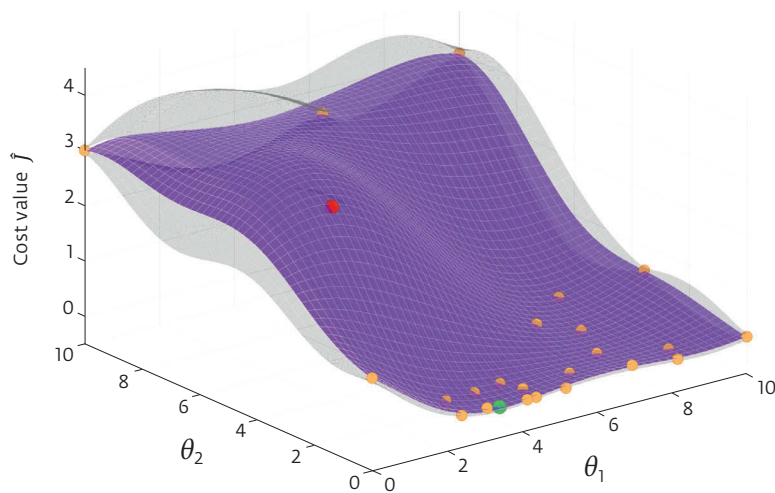
Affordable control: The mathematical cost function describes how well a robot balances a rod; it does so particularly well when the cost value is low. The cost function depends on the parameters θ_1 and θ_2 of the feedback control algorithm the machine uses to control its movements. When Apollo balances a long rod on his hand, the controller that worked well for the short rod (red dot) proves unsuitable. With the help of a learning algorithm, Apollo then systematically tries out new controllers (yellow dots) until he finds the one best suited for solving the new task (green dot).

The computer scientist has built up a database into which she entered models of more than 700 objects – from hammer to toy doll. In order for the robot to learn how to successfully grasp these objects, she simulates countless possible gripping techniques on her computer. In the process, she also takes into account that an object's position might shift if the robot touches it with its fingertips first instead of with its palm, which might move the object as the robot tries to grasp it. That might just cause the object to slip into the robot's hand, or it might not.

The aim is that one day, based on the experience the software gains from these simulations, robots will be able to grasp not only things they were taught to recognize, but also unfamiliar objects – even if their sensors provide them with only incomplete or noisy information.

Helping robots grab hold of things is also the goal of Ludovic Righetti, a Research Group Leader at the Max Planck Institute in Tübingen. While Jeannette Bohg works on teaching robots to use visual information to develop a plan for grasping an unknown object, Ludovic Righetti and his team approach this challenge from a different angle: among other things, they teach students like Apollo to grasp objects more sensitively. The aim is for a robot hand to be able to grab hold of

an object even if the hand doesn't make contact at the right points.

Such actions are regulated by a feedback control system – a computer program that creates a feedback loop between the information collected by the sensors and the motions performed by the actuators. In Apollo's case, the feeling of "I've got it" or "I didn't get it" is expressed as data measured by the force sensors in his hand. The control unit in his brain translates this data into a command directed at the actuators in his fingers. The sensors then report whether the fingers really did end up in the planned location. If not, the software corrects the robot's motions. This type of control engineering is always based on a model that expresses the design of a robot and the interaction between its control unit and actuators as mathematical formulas.

FEEDBACK CONTROL SOFTWARE LEARNS INDEPENDENTLY

Developing the correct model for a tin man or woman is in fact a highly complex affair: "The physics of a robot are extremely nonlinear," says Stefan Schaal. In other words, small deviations from the model's assumptions, for example with regard to sensor sensitivity or actuator force, can have dire consequences. The robot might go complete-

ly haywire; in any case, it won't do what it is supposed to do. The main reason is because a full-body robot has around 40 degrees of freedom: it can move its various limbs with the help of 40 independent joints.

But the actual problem doesn't even lie in the physical model on which the robot's controls are based; the model can be controlled despite any adverse circumstances. "I can develop a good model of my robot, but not of unfamiliar surroundings," explains Schaal.

That's why part of Righetti's team is using machine learning to teach robots to develop a more flexible model for solving a particular task, such as grabbing hold of a cup. "This way, the robot learns how an action is supposed to feel at any given point in time – that is, what the force sensors in its wrists, the haptic sensors in its fingers, and the camera eyes are supposed to be registering," says Righetti. "This is a relatively simple form of learning." If the grip turns out to be wrong, Apollo and his fellow students can correct it using their adaptive controls. "Ultimately, we hope to develop more general models that can be applied for a wide range of tasks."

The approach used by Righetti's team involves models that know, or at least should know, which actuator force leads to which movement. The researchers then control the exerted force

» Robotics must not follow just one single path when developing autonomous machines designed to serve as household assistants or emergency response units.

and thus the actions performed by the machine. Most robotics scientists, in contrast, currently employ control systems whose commands explicitly specify which position a robot's hand or foot is supposed to assume.

While this might sound like a mere technical detail, it impacts the overall implementation: if control systems measure their success in terms of whether or not a hand reached its target position, for example, the hand won't let anything – not even a human – stop it from carrying out its orders. If need be, the robot will apply even more force to assert itself. In the case of industrial robots, which typically use particularly high levels of force, this could lead to serious accidents.

If, however, the exerted force is regulated, the machine can be programmed to be more sensitive – an indispensable prerequisite for interacting with humans. "Our force-controlled robots are able to act in a much more compliant manner without compromising their precision, because we use controllers that are robust to imprecise models," says Righetti. "This approach opens up a range of new possibilities and is sure to become more widely adopted in the future."

Another researcher working on improving feedback control in robots with the help of machine learning is Sebastian Trimpe. You could say he develops the class materials that help robots learn how to balance a rod, for example, much like children learn to balance a stick on just one finger. "That is a relatively simple task," says Trimpe. "But once we understand how a robot best learns how to solve it, we may also be able to teach it to learn more sophisticated skills." For instance, standing and walking on unfamiliar and uneven terrain.

ROBOT CURIOSITY IS GUIDED BY METHODS OF PROOF

Apollo can balance a rod thanks to his internal feedback control algorithm, which analyzes the sensor information that indicates the rod's current position and movement, and translates this data into control signals directed at the actuators. So if the rod is about to tilt to the right, for example, the controller intervenes and corrects Apollo's movement to prevent the rod from tipping over.

In fact, Apollo's teacher even makes the task a bit more challenging by having him first balance a shorter rod. It's more difficult to balance a shorter rod than a longer one, because the short rod has less inertia and is therefore more likely to tip over sooner, making quick corrective action necessary. But Apollo balances the short rod with ease, even though his first attempts at doing the same with the longer rod fail miserably.

Sebastian Trimpe isn't surprised by the unsuccessful attempt: The corrective action needed depends on the length of the rod, which the controller takes into account. However, the researchers hadn't yet adjusted the control algorithm when they handed Apollo the longer rod. This means the feedback control that worked fine for the short rod fails for the long rod, as it causes Apollo to move his arm much too abruptly.

"Instead of programming a new algorithm for each new rod, we adjusted the control software to allow it to learn independently," says Sebastian Trimpe. With the help of machine learning, the robot can therefore autonomously adapt to a new situation without this being preprogrammed into its system. In control engineering, a domain of classical engineering, this is a relatively new approach.

Furthermore, the researchers have programmed instructions that train Apollo to independently learn the best controller in as few attempts as possible. "The algorithm automatically suggests the controller that offers the great-

A sure footing: Ludovic Righetti's team has optimized Hermes' controls so that the robot quickly corrects its posture when it has lost its balance. This prevents the robot from falling down even if it is pushed or is standing on shaky ground.
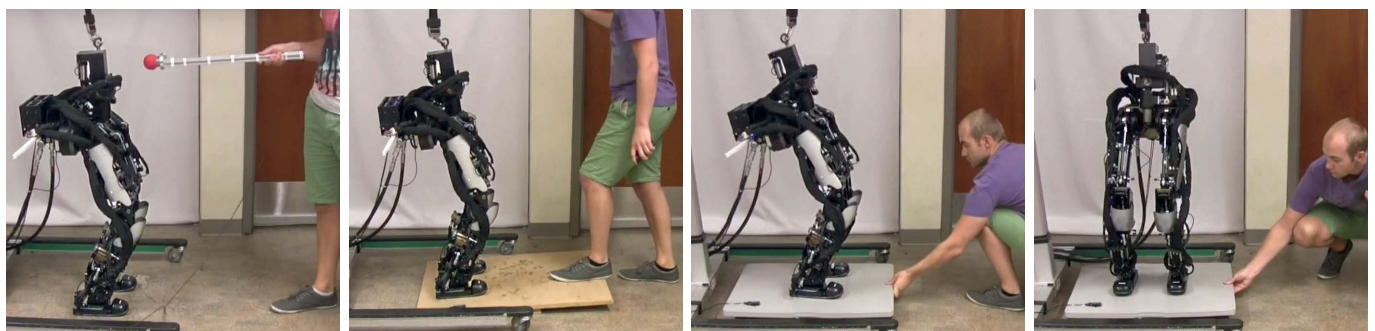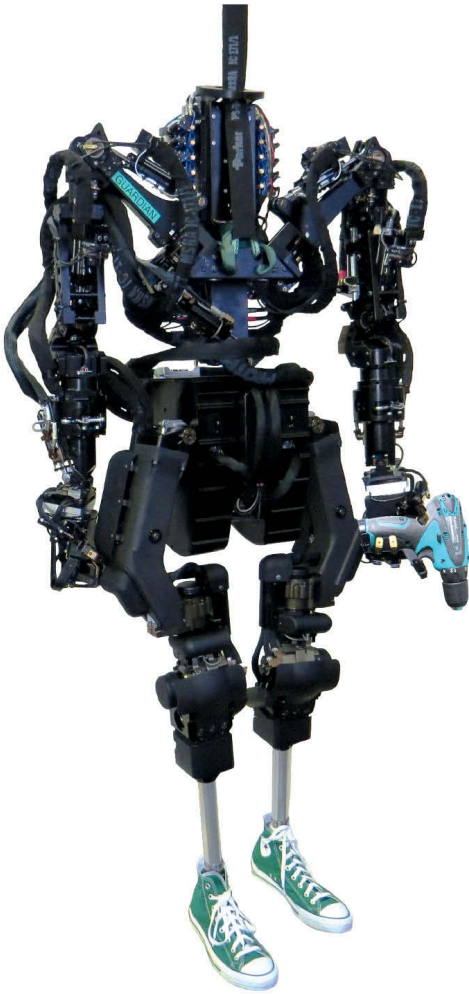
A Goddess in Chucks: Athena uses hydraulics to move her limbs when handling a tool or walking. She may not have a head, but she's wearing trendy sneakers.

est learning effect," explains Trimpe. In the early stages of the learning process, these could be controllers that differ significantly from the original controller. That's why Apollo's second and third attempts at balancing the rod appear even clumsier than the first, which doesn't faze the robot. After that, however, the learning curve rises steeply.

"In contrast to typical machine learning applications, such as image recognition, for example, learning in robotics is a dynamic problem," says Trimpe. The data set the software uses to recognize faces doesn't change. A robot, on the other hand, constantly gathers new information and gains new experience as it moves through and interacts with the world around it. That's why learning should be a lifelong process for a robot. Yet that very goal puts it in a predicament time and again.

"In order to learn new things or improve, the robot must try out new behaviors," says Trimpe. That can also mean that it performs more poorly for a period of time. To prevent the robot from getting up to nonsense or even becoming damaged in the event of a fall, the researchers must integrate guarantees into the learning algorithm. Using mathematical proof techniques, they guide the robot's curiosity to ensure that the behavior it learns is not only flexible, but also sensible and robust.

In an effort to make a robot's controls more robust, meaning less prone to failure, Ludovic Righetti focuses on more than just the machines' ability to learn. His work is a prime example of modern robotics not pursuing just one single path in developing machines that could one day serve as domestic help or emergency response units.

POSTURE CORRECTED IN JUST A FEW MILLISECONDS

"We want to take the approach in which a robot develops models based on experience, and combine it with a different control engineering approach," says Righetti. He and his team program the flexibility needed to quickly correct the robot's posture directly into the algorithms of the control unit that creates the feedback loop between the sensor data and the commands for the actuators. Whether or not this results in a sensible action, such as grabbing hold of a cup or balancing on shaky ground, can, from a control engineering point of view, be formulated as a mathematical optimization problem, the solution to which identifies the most suitable controller for the particular task at hand.

A controller can often be optimized before the robot is put into operation. When that is not the case due to unforeseen circumstances – for example, if the robot trips or is pushed – the controller must be improved mid-action, for example while the robot is walking. "We have developed strong algorithms for this purpose," says Ludovic Righetti.

Not only do the methods reliably compute how the controller must be adjusted to accommodate for unexpected events, but the software is also very fast – an absolute must, especially if the robot is to walk across uneven terrain. "In that type of situation, the robot has only a few milliseconds to correct its posture once it starts losing its balance," says Ludovic Righetti. If the robot fails to do so, it will fall to the ground.

Righetti and his team use Hermes as proof of how well a machine is capable of using its control system to maintain its balance. In a way, Hermes is Apollo's counterpart, as he consists of only a lower body and two legs. And there is a good reason why he has only two legs: while a robot does in fact have a more secure footing with four or more legs or even wheels, there are many obstacles that it can overcome only by climbing over them using two legs and two arms.

Stefan Schaal and Ludovic Righetti experiment with the two-legged Hermes at the University of Southern California, where they both conducted research before joining the Max Planck Institute in Tübingen. When the researchers throw Hermes out of balance by pushing him, for example, he corrects his posture using distinctly human-like movements.

The new control system is catching on: "The same techniques are now being used with many robots," says Righetti. The mechanism will also help the newest addition to the Tübingen-based institute's Mount Olympus maintain her balance: Athena, the first robot to fly from the US to Germany while sitting in a normal passenger seat. She,

Robotics necessitates mathematics: Ludovic Righetti (left) and Sebastian Trimpe discuss the control algorithms that help Apollo and his fellow students learn to act autonomously and flexibly.

too, doesn't quite live up to the expectations associated with the goddess after whom she was named. Her bulky torso, strong arms, hydraulics tubes and machine head, which she doesn't even always wear, make her about as graceful as a transformer. But at least she has both arms and legs.

The research conducted by Ludovic Righetti and his team will now focus on identifying the best way for Athena to coordinate her limbs while attempting to solve several tasks simultaneously. They also seek to answer the question as to how she sets the right priorities when trying to grasp an object while walking or standing on uneven ground. After all, the machine initially doesn't know that it's more important to stay upright than to unconditionally try to grab the object in front of it. A robotic assistant is suitable for day-to-day tasks only if it is capable of making sensible decisions.

It will take some time before robots are independent enough to be able, for instance, to help people in need, as the field of robotics still needs to make significant adjustments to many points in the perception-learning-action loop.

That's why Stefan Schaal doesn't believe that, 30 years from now, we will be cared for by beings made of tin, plastic and electronic components – not only owing to technological issues, but also because of society's possible reluctance to accept it. "But I hope I live to see the day when robots help us with our domestic chores, even if only by picking up a book when told to do so because we can't bend down anymore," says the scientist.   ◄

TO THE POINT

- **Modern-day robots aren't yet capable of flexibly adapting to new tasks and unexpected situations. Moreover, they are prone to error.**
- **Using machine learning and other methods, Max Planck researchers in Tübingen aim to teach robots to perform such tasks as quickly and reliably finding objects in unfamiliar settings, securely grasping previously unknown objects and independently learning the most suitable control system for solving new tasks.**
- **In order to help robots gain a surer footing and a more stable gait, the researchers also program the controls in such a way that the machines continuously optimize their actions and thus react to disruptions or unforeseen events.**

GLOSSARY

Machine learning: Using large amounts of data, a software program learns examples of a particular type of task and can subsequently carry out this task in a general manner. A large number of pictures in which faces are specially indicated, for example, tell a software program which features are essential for facial recognition. The program is then able to identify faces on pictures it has never seen before.

Feedback control: When a machine uses data collected by sensors to issue a command for a certain action to be performed, and when it uses this sensor data to monitor the execution of the command and, if necessary, correct the action being carried out, this is known as a feedback control system. Open-loop control systems don't include this feedback function.