


Robot Discovers Itself, Adapts to Injury



Graduate student Viktor Zykov, former student Josh Bongard, now a professor at the University of Vermont, and Hod Lipson, Cornell assistant professor of mechanical and aerospace engineering, watch as a starfish-like robot pulls itself forward, using a gait it developed for itself. The robot's ability to figure out how it is put together, and from that to learn to walk, enables it to adapt and find a new gait when it is damaged. Credit: Lindsay France/Cornell University

Nothing can possibly go wrong ... go wrong ... go wrong ... The truth behind the old joke is that most robots are programmed with a fairly rigid "model" of what they and the world around them are like. If a robot is damaged or its environment changes unexpectedly, it can't adapt.

So Cornell researchers have built a robot that works out its own model of itself and can revise the model to adapt to injury. First, it teaches itself to walk. Then, when damaged, it teaches itself to limp.

Although the test robot is a simple four-legged device, the researchers say the underlying algorithm could be used to build more complex robots that can deal with uncertain situations, like space exploration, and may help in understanding human and animal behavior.

The research, reported in the latest issue (Nov. 17) of the journal *Science*, is by Josh Bongard, a former Cornell postdoctoral researcher now on the faculty at the University of Vermont, Cornell graduate student Viktor Zykov and Hod Lipson, Cornell assistant professor of mechanical and aerospace engineering.

Instead of giving the robot a rigid set of instructions, the researchers let it discover its own nature and work out how to control itself, a process that seems to resemble the way human and animal babies discover and manipulate their bodies. The ability to build this "self-model" is what makes it able to adapt to injury.

"Most robots have a fixed model laboriously designed by human engineers," Lipson explained. "We showed, for the first time, how the model can emerge within the robot. It makes robots adaptive at a new level, because they can be given a task without requiring a model. It opens the door to a new level of machine cognition and sheds light on the age-old question of machine consciousness, which is all about internal models."

The robot, which looks like a four-armed starfish, starts out knowing only what its parts are, not how they are arranged or how to use them to fulfill its prime directive to move forward. To find out, it applies what amounts to the scientific method: theory followed by experiment followed by refined theory.

It begins by building a series of computer models of how its parts might be arranged, at first just putting them together in random arrangements. Then it develops commands it might send to its motors to test the models. A key step, the researchers said, is that it selects the commands most likely to produce different results depending on which model is correct. It executes the commands and revises its models based on the results. It repeats this cycle 15 times, then attempts to move forward.

"The machine does not have a single model of itself -- it has many, simultaneous, competing, different, candidate models. The models compete over which can best explain the past experiences of the robot,"

Lipson said.

The result is usually an ungainly but functional gait; the most effective so far is a sort of inchworm motion in which the robot alternately moves its legs and body forward.

Once the robot reaches that point, the experimenters remove part of one leg. When the robot can't move forward, it again builds and tests 16 simulations to develop a new gait.

The researchers limited the robot to 16 test cycles with space exploration in mind. "You don't want a robot on Mars thrashing around in the sand too much and possibly causing more damage," Bongard explained.

The underlying algorithm, the researchers said, could be applied to much more complex machines and also could allow robots to adapt to changes in environment and repair themselves by replacing parts. The work also could have other applications in computing and could lead to better understanding of animal cognition. In a way, Bongard said, the robot is "conscious" on a primitive level, because it thinks to itself, "What would happen if I do this?"

"Whether humans or animals are conscious in a similar way -- do we also think in terms of a self-image, and rehearse actions in our head before trying them out -- is still an open question," he said.

Source: Cornell University

This document is subject to copyright. Apart from any fair dealing for the purpose of private study, research, no part may be reproduced without the written permission. The content is provided for information purposes only.