THE UNIVERSITY OF VERMONT

Release Date: 11-17-2006

'Science' Features Professor's Resilient Robot

Contact: University Communications Staff
Email: newserv@uvm.edu
Phone: (802) 656-2005 **FAX:** (802) 656-3203

A University of Vermont computer scientist helped develop the first robot capable of detecting its own shape and using this knowledge to efficiently adapt to damage. The new technology may have applications for robots used as planetary rovers or in disaster relief — it also has philosophical implications, shedding light on artificial consciousness and curiosity.

Joshua Bongard, an assistant professor in the College of Engineering and Mathematical Sciences, is lead author of a paper describing the project, "Resilient Machines Through Continuous Self-Modeling," that appeared in the Nov. 17 issue of the journal *Science.* The other co-authors and project participants are Victor Zykov and Hod Lipson of Cornell University.

The group's robot, which resembles a starfish, initially has no sense of its own shape. It measures the results of a limited number of small movements to develop plausible models of its shape and construction. The robot evaluates and refines these competing candidate models through more movements and observation, and efficiently arriving at an accurate internal model of its shape. The robot can then use this continuously updated self-model to detect damage and develop new ways to move even after sustaining damage like the loss of a leg.

More traditional robots either require extensive programming to adapt to damage, dramatically limiting their adaptability in the face of unexpected conditions, or they must experiment far more extensively and mechanically to develop new ways of moving, limiting efficiency and increasing the risk of damage in hostile environments.

Bongard answers some questions about the robot below.

What was significant about this project?

JOSHUA BONGARD: The most important thing here for us is this is the first robot that can build up a description of its own body. So the robot can build up a sense of self; that hasn't been done before in robotics. The second interesting thing about this is that it then uses that self-model, that sense of self, to actually try out different ways of moving. We commanded this robot to *learn* how to move; we didn't tell the robot how to move. It tries internally using this self-model, "What would happen if I tried hopping? What would happen if I crawled?" And so on. And eventually it comes up with a behavior that it thinks will actually work and then tries it out in reality; more often than not the robot starts moving.

How does that approach contrast with more traditional ideas how to control a robot?

There are two existing approaches. In the first, the idea is to allow the robot to attempt hundreds or thousands of trials in the real world, and eventually it hits on a way of moving. In our case, we're dealing with a robot that is damaged. Potentially, for example, this could be a robot probe on a remote planet, and we don't want it thrashing around wildly because it might damage itself further or fall off a cliff. We want it to be very careful about what it does, and perform as few exploratory trials as possible. The second existing approach is to create by hand a model for the robot. The roboticist would tell the robot, you're made up of four legs, and you're put together in this way, and you can do this and you can't do that. That approach severely limits the intelligence or the adaptability of the robot. The robot in that situation can't very easily adapt and overcome unanticipated situations.

Where does this go from here?

We basically developed this as a proof of concept for ideas for the next generation of planetary rovers. NASA is very interested in having a robot like this… We can't assume that the robot can easily communicate back with mission control on Earth and communicate what it's sensing and what it should do next. We want the robot to figure out on its own how it should go about exploring the surface of the planet. The other application would be for deploying these robots in a disaster site. A disaster site, like the surface of a remote planet, is a very unpredictable environment and there's a high likelihood that the robot may become damaged, so again we want the robot to quickly adapt and carry on with its mission.

How does your part in this fit into your larger intellectual interests?

There's a practical interest here, but for myself in particular, what's more interesting is the conceptual side of things. This robot starts to suggest something about the nature of curiosity, in the sense that the robot, when it's learning about itself, doesn't simply thrash around randomly. It actually tries out each time a new action to try to learn something new about its own body and its local environment. In a sense, at a very rudimentary level, this robot is curious.

It also suggests something about the nature of self-awareness. This robot starts by having little awareness of its own body, and through interaction through the physical world it gains experience and builds up a sense of itself, a simulation of its own body, and it can then come to understand what that body is capable of and what it isn't capable of. Taking that a step forward then, perhaps we can start someday to use robots as tools to start to ask questions about the nature of human self-awareness and curiosity. Is there something going on in our brains similar to what's going on in the brain of this robot?
