DISCOVER
Science, Technology, and The Future

- [Subscribe](#)
- [Give a Gift](#)
- **[Archives](#)**
- **[Customer Service](#)**

SEARCH [   ] [SUBMIT»]

- [Health & Medicine](#)
- [Mind & Brain](#)
- [Technology](#)
- [Space](#)
- [Human Origins](#)
- [Living World](#)
- [Environment](#)
- [Physics & Math](#)
- [Video](#)
- [Photos](#)
- [Podcast](#)
- [RSS](#)

[Blogs](#) / [Science Not Fiction](#)

« [How to Be A More Human Human](#)

Robots That Evolve Like Animals Are Tough and Smart—Like Animals

🔲 📘 📄   ⁔

0   digg   🌀 **Stumble!** **Like?**

People who work in robotics prefer not to highlight a reality of our work: robots are not very reliable. They break, *all the time*. This applies to all research robots, which typically flake out just as you're giving an important demo to a funding agency or someone you're trying to impress. My fish robot is back in the shop, again, after a few of its very rigid and very thin fin rays broke. Industrial robots, such as those you see on car assembly lines, can only do better by operating in extremely predictable, structured environments, doing the same thing over and over again. Home robots? If you buy a Roomba, be prepared to adjust your floor plan so that it doesn't get stuck.

What's going on? The world is constantly throwing curveballs at robots that weren't anticipated by the designers. In a novel approach to this problem, Josh Bongard has recently shown how we can use the principles of evolution to make a robot's "nervous system"—I'll call it the robot's controller—robust against many kinds of change. This study was done using large amounts of computer simulation time (it would have taken 50–100 years on a single computer), running a program that can simulate the effects of real-world physics on robots.

What he showed is that if we force a robot's controller to work across widely varying robot body shapes, the robot can learn faster, and be more resistant to knocks that might leave your home robot a smoking pile of motors and silicon. It's a remarkable result, one that offers a compelling illustration of why intelligence, in the broad sense of adaptively coping with the world, is about more than just what's above your shoulders. How did the study show it?

Each (simulated) robot starts with a very basic body plan (like a snake), a controller (consisting of a neural network that is randomly connected with random strengths), and a sensor for light. Additional sensors report the position of body segments, the orientation of the body, and ground contact sensors for limbs, if the body plan has them. The task is to bring the body over to the light source, 20 meters away.

A bunch of these robots are simulated, and those that do poorly are eliminated, a kind of *in-computo* natural selection. The eliminated robots are replaced with versions of the ones that succeeded, after random tweaks ("mutations") to these better controllers have been made. The process repeats until a robot that can get to the light is found. So far, there's been no change in the shape of the body.

With the first successful robot-controller combination found (one that gets to the light), the body form changes from snake-like to something like a salamander, with short legs sticking out of the body. (All body shape changes are pre-programmed, rather than evolved.) The evolutionary process to find a successful controller-bot combination repeats, with random changes to the better controllers until, once again, a controller-bot combination is found that is able to claw its way to the light.

Then the short legs sticking out to the side slowly get longer, and rather than sticking out to the side, they progressively become more vertical. With each change in body shape, the evolutionary process to find a controller repeats. Eventually, the

sim-bot evolves to something that looks like any four-legged animal.

That was all for round one of evolution. For round two, the best controller from round one was copied into the same starting snake-like body type that round one began with. But now, the change in body forms occurs more rapidly, so that by the time 2/3 of the "lifetime" of the robot is completed, it has reached its final dog-like form. For round three, this all happens within 1/3 of the robot's lifetime. For round four, the body form starts off as dog-like and stays there.

So there are changes occurring at two different time scales: changes over the "lifetime" of the robot, similar to our own shape changes from fetus to adulthood; and changes that occur over generations, through which development during a lifetime occurs more rapidly. The short time scale is called "ontogenetic" and the long scale (between the different rounds) is "phylogenetic."

The breakthrough of the work is that it found that having these variations in body shape occur over ontogenetic and phylogenetic time scales resulted in finding a controller that got the body over to the light much faster than if no such changes in body shape occurred. For example, when the system began with the final body type, the dog-like shape, it took much longer to evolve a solution than when the body shapes progressed from snake-like to salamander to dog-like. Not only was a controller evolved more rapidly, but the final solution was much more robust to being pushed and nudged.

The complexity of the interactions over 100 CPU years of simulated evolution makes the final evolved result difficult to untangle. Nonetheless, there is good evidence that the cause of accelerated learning in the shape-changing robots is that the controllers developed through changing bodies have gone through a set of "training-wheel" body shapes: a robot starting with a four-legged body plan and a simple controller quickly fails—it can't control the legs well and simply tips over. Starting with something on the ground that slithers, as was the case in these simulations, is less prone to such failures. So not any old sequence of shape changes works: mimicking the sequence seen in evolution garners some of the advantages that presumably made this sequence actually happen in nature, such as higher mechanical stability of more ancient forms.

Less clear is the source of increased robustness—the ability to recover from being nudged and pushed in random ways. Bongard suggests that the increased robustness of controllers that have evolved with changing body shapes is due to those controllers having had to work under a wider range of sensor-motor relationships than the ones that evolved with no change in body shape. For example, any controller that's particularly sensitive to a certain relationship between, say, a sensor that reports foot position, and one that reports spine position would fail (and thus be eliminated) as those relationships are systematically changed in shifting from salamander-like to dog-like body form and movement. So that means that if I suddenly pushed down the back of a four-legged dog-like robot, so that its legs would splay out and it would be forced to move more like a salamander, the winners of the evolutionary competition would still be able to work because the controllers had worked in salamander-like bodies as well as in dog-like bodies.

In support of this idea, the early controllers, that were purely based on moving the body axis ("spine"), appear to be still embedded in the more advanced controllers; so if something happens to the body (say, one leg gets knocked), the robot can revert to more basic spine-based motion patterns that don't require precise limb control. Bongard observed that the controllers evolved through changing body shape exhibited more dependence on spinal movement, using the legs more for balance, than those evolved without changing body shape. (It would be interesting to try his approach with simulated aquatic robots, which can be neutrally buoyant like many aquatic animals are, and thus don't have the "tipping over" problem that Bongard's simulated terrestrial robots had).

To be fair to existing robots, even with a controller that worked under every conceivable body shape and environmental condition, they would still break all the time. This is because the materials we make them out of are not self-healing, in contrast to the biomaterials of animals. Animals are also constantly breaking (at least on a micro level), and the body constantly repairs this. Bones subjected to higher loads, like the racket arm of a tennis player, get measurably thicker. Not only is the body self-repairing, recent innovative computer simulations of real neurons that generate basic rhythms like walking and chewing have shown that the neurons keep generating the rhythm despite big variations in the functioning and connections of these neurons. These functions are so important to continued existence—the body's version of too big to fail—that embedded within them are solutions to just about everything the world can throw at them.

This new work provides the fascinating and useful result that fashioning controllers that work through a sequence of body shapes mimicking those seen in evolution accelerates the learning of new movement tasks and increases robustness to all the hard knocks that life inevitably delivers. It suggests that without the sequence of body shapes that evolution and development bring about, we might have nervous systems that are much too finely tuned to our adult upright bipedal form. Instead of crawling to help after we twist our ankle in the woods, we'd be left with nothing but howling for help.

February 14th, 2011 Tags: embodiment, evolution
by Malcolm MacIver in Artificial Intelligence, Biology, Computers, Cyborgs, Genetics, Neuroscience, Robots | 1 comments | RSS feed | Trackback >

One Response to "Robots That Evolve Like Animals Are Tough and Smart—Like Animals"

1. 1. *Douglas Watts* Says:
   February 14th, 2011 at 11:36 pm

   Can you make a passenger pigeon or an ivory billed woodpecker or a sea mink or a Stellers sea cow?

Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Submit Comment

- ## About Science Not Fiction

  Sometime in the future, a group of renegade scientists and technologists will take a time machine to now. They're spilling the secrets of tomorrow here at Discover's Science Not Fiction blog.

  ▪ **Kevin Grazier** is a planetary physicist who works at NASA/JPL on the Cassini/Huygens Mission to Saturn and Titan. He also teaches at UCLA and Santa Monica College, and is a science advisor for sci-fi shows like *Eureka* and *Battlestar Galactica*. Kevin is co-author of the upcoming book "The Science of Battlestar Galactica," and covers space and physics for SNF.

▪ **Malcolm MacIver** is a bioengineer at Northwestern University who studies the neural and biomechanical basis of animal intelligence. He consults for sci-fi films (*Tron Legacy*, Joss Whedon's *The Avengers*), and was the science advisor for *Caprica*. He covers AI and robotics for Science Not Fiction.

▪ **Kyle Munkittrick (Web, Twitter)** is program director at the Institute for Ethics and Emerging Technologies and a grad student at NYU. He covers transhumanism.

- # Subscribe

The RSS feed for Science Not Fiction is here .

- **80beats**

Categories

Copyright © 2010, Kalmbach Publishing Co.

Privacy - Terms - Reader Services - Subscribe Today - Advertise - About Us