

Robots are learning how to limp like animals

An algorithm can help bots adapt to injury through trial and error

By [Adi Robertson](#) on May 27, 2015 01:00 pm

Robots somehow seem most lifelike when they're not working quite right. Watching a spidery mechanical creature struggle to keep going after losing a limb feels like watching a wounded animal, with one big difference: the animal has had a whole life to figure out how to adapt to an injury, not to mention millions of years of evolution. The robot might not have any experience at all. That's the problem that Sorbonne University researcher Jean-Baptiste Mouret and a team of fellow scientists are trying to solve. In this week's issue of *Nature*, the group lays out an algorithm that could potentially help robots compensate for injury faster and more consistently.

"It's quite intuitive for us humans to say 'if something is broken, first I want some kind of diagnosis to know what's wrong, and then we will try to find a solution for it,'" says Mouret. Sometimes this takes the form of set patterns written by humans, and sometimes robots can figure it out on their own. Researchers at Cornell University, for example, developed a [four-legged "starfish" robot](#) that could identify parts of its own body and adapt to changes accordingly. Mouret's team, though, thinks robots could adapt simply through trial and error, without having to understand what's wrong first — the way someone might start limping without knowing whether they've got a sprain or a fracture. And they think that by "teaching" the robot about what its own body can do first, they can help it draw on past knowledge to do so.

"IT DOESN'T UNDERSTAND WHAT IS BROKEN. IT JUST KNOWS THAT SOMETHING IS NOT AS EXPECTED."

For this study, once the team built and programmed a robot, the algorithm catalogued its potential range of behavior — every way that the robot could conceivably move — and identified the best options. It's "a bit like a child who is trying many different things with its own body, just to know itself," says Mouret. By the time it was sent into the field, the robot had a set of basic intuitions based on its undamaged body. After getting hurt, it could use that knowledge as a jumping-off point to test new options. "These intuitions are wrong most of the time," says Mouret. "The robot does not trust them, but says 'let's try this thing, because we have intuition that that might work.'"

Once it's started this process, the robot uses a mathematical model called Bayesian optimization to find a new strategy. A robot, for example, might begin by realizing that its normal walking style is now sending it in circles — researchers have deactivated one of the legs. It checks its existing knowledge base of around 13,000 gaits, picks a likely solution, and tests it, while the sensors measure its speed and direction. Then it moves on to the next option, balancing its "intuition" with more random choices that could reveal a new path. Once it's 90 percent confident that it's found the fastest and straightest option, it gets back to work. "It doesn't understand what is broken. It just knows that something is not as expected and then tries to find something that works in spite of the damage," says Mouret.

The team tested this algorithm on two robots: a six-legged walking bot and a multi-jointed, rotating arm that was attempting to drop a ping-pong ball into a cup. They tested five different ways of breaking the hexapod, including partially severing one leg, completely removing two, and attempting a makeshift repair. After being run through each scenario 40 times, the it generally took about a

minute to adapt to a new gait. The robotic arm, meanwhile, had individual joints either broken or stuck in place; it adapted consistently in under two minutes and usually under 30 seconds.

A damaged hexapod (Antoine Cully / UPMC).

The model has limits, though. For two of the 14 mechanical arm errors they tested, for example, the team's algorithm didn't perform any better than traditional, non-"intuitive" trial and error. The team hypothesizes that if the damage is bad enough, the robot's body can change so much that its previous knowledge becomes useless. Likewise, if a robot's sensors — not just its appendages — are damaged, it won't be able to get enough accurate information to adapt at all. And right now, it only responds to fairly simple goals. When humans and animals limp, for example, they're not just trying to find a new way to walk, they're avoiding hurting themselves any further. There's no equivalent to that for robots; they're just looking to be faster or more accurate.

"It's very difficult to imagine how blind the robots are," says Mouret. "We have skin, we have all of our internal organs, all of these things are sending sensory information. But robots usually have just a few sensors." Partly because of this, the robots Mouret and his team are testing also barely know anything about what's going on outside their own bodies. So if they're on uneven ground, or in a small space, they can't factor that into their trial and error. The next step is to rectify this, testing more complicated robots with a wider range of motion and more ways of looking at the world.

"IT'S VERY DIFFICULT TO IMAGINE HOW BLIND THE ROBOTS ARE."

Overall, Mouret believes that this system could be adapted easily to existing

robots, particularly advanced humanoid ones like the competitors in DARPA's [Robotics Challenge](#). "All these robots are very versatile," he says. If one part breaks, they've got a lot of ways to compensate, just like humans or animals. The more stripped-down and specific a machine gets, the less useful the algorithm becomes. "If you take a Roomba — the vacuum cleaner — this robot has only two wheels. So if one wheel is broken, maybe it can do something. If two are broken, it cannot do anything," says Mouret. "So it could be used, but it would not be that interesting." Of course, it can be just as important for fairly inflexible systems like jet aircraft to respond to damage. Researchers at NASA, for example, are testing a [flight control system](#) with its own intelligent adaptive system.

In theory, helping robots figure out what they can do and then telling them to achieve a specific goal could have applications beyond adapting to injuries. An undamaged walking robot, for example, could use the algorithm to figure out how to traverse difficult terrain. But these are secondary options. "I think engineers are very good at their jobs. So when we try to optimize something and make it work or make a better robot, they usually have very good ideas," says Mouret. Engineers just can't predict everything that might go wrong with their creations. "This kind of algorithm could be used to optimize robots and to make them perform better. But I think it's much more useful when nobody is around."

This is Jared Leto as
the Joker. I'm sorry.

This is Harley Quinn
and the rest of the
Suicide Squad

NASA just released
the best GIFs of Pluto
so far

[Learn more](#)

Powered by

for you