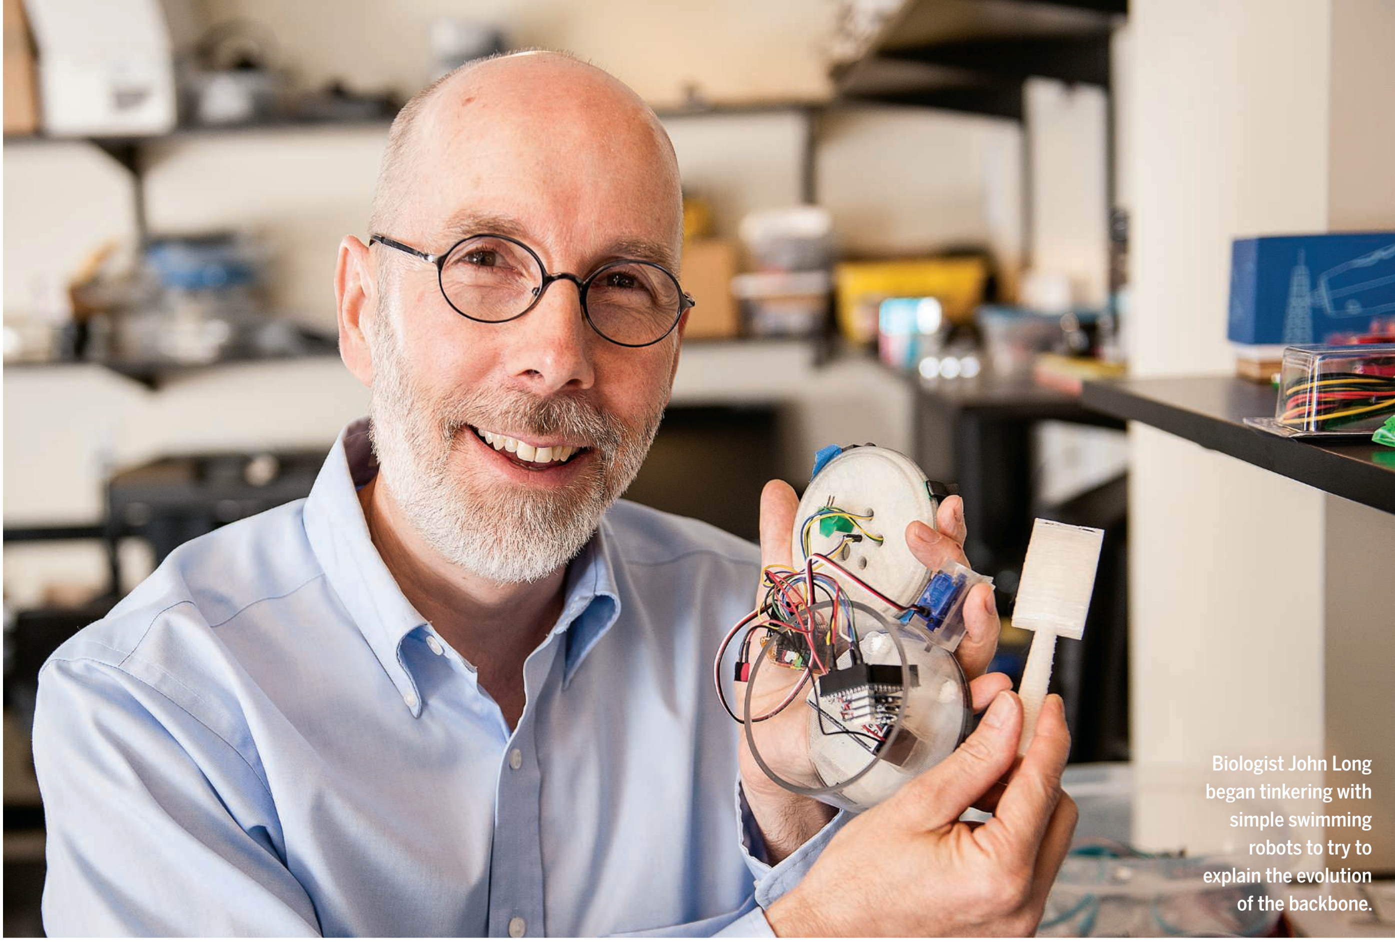

Biologist John Long began tinkering with simple swimming robots to try to explain the evolution of the backbone.

The accidental roboticist

John Long wondered how life developed the capacity to evolve—so he unloosed a fleet of robot tadpoles

By **Adrian Cho**, *in Poughkeepsie, New York*

The Tadro is a primitive creation. The robotic tadpole's palm-sized cylindrical body contains a few wires and resistors and a reprogrammable chip. An electric motor wags the tail, a simple plastic cone tapering to a square fin. On either side of the robot's head perch two light sensors, with a third in the center. That's it. Yet when placed in its environment—for now, a kiddie pool in a darkened locker room here at Vassar College—the Tadro does something that most complex machines cannot: Unlike your car or computer, it guides itself and behaves like a living creature.

A flood lamp hangs a meter or so over the pool, supplying light that serves as Tadro's "food." Tipping slightly forward, the robot churns toward the lamp, its motor squeaking *weekee weekee weekee*, and then circles under it, feeding on the glow. As the Tadro wriggles toward the light again and again, it's hard *not* to think it's alive.

And like living things, Tadros evolve. With help from their makers, they change from generation to generation in response to a form of survival-of-the-fittest selection. They are the brainchildren of John Long, a Vassar biologist with a cheerful smile and a scholar's little round glasses who peppers his conversation with references to books and movies. ("Pay no attention to the man behind the curtain!") He has used Tadros to study the evolution of backbones, testing the idea that by making ancient fish stiffer, backbones made them faster and hence better at collecting food or evading predators.

Now, he and his team are gearing up for an even more ambitious effort. They plan to use Tadros to probe the font of all life's diversity: the ability to evolve, or evolvability. One key to that ability, he and his collaborators think, may be modular design, especially in the brain. In animals, distinct neural circuits control different functions, such as vocalization and vision. "The grand hypothesis is that modularity will enhance evolvability," Long says. "It's this capacity for future change that we're trying to get our hands around." If he is right, modularity itself should evolve within the Tadro's control circuits, under the right conditions.

That may be a lot to ask of toy tadpoles, but others say that experiments with robots can lay bare the nuts and bolts of evolution in ways that observations with living things cannot. "You're able to set up and test hypotheses that couldn't be tested otherwise," says Robert Pennock, a philosopher of science and evolutionary biologist at Michigan State University in East Lansing. A. E. "Gusz" Eiben, an evolutionary computer scientist at VU University Amsterdam, agrees. Long's work, he says, "deserves to be more widely known, especially among biologists."

IN THE YOUNG FIELD of evolutionary robotics, Long's research swims against the prevailing current. Most researchers use evolution as a tool to develop better robots. They construct otherwise identical robots with a variable trait—say, the length of a limb or some aspect of the robot's circuitry—which is specified in an abstract numerical "gene." They set the robots loose in some environment to determine, according to some previously decided criterion related to their behavior, which ones are fitter and get to pass their winning traits to more offspring.

There's no hot sex for the robots, however. Instead, mating occurs entirely within a computer using a "genetic algorithm" to determine the traits of the next generation of robots. To mimic biological reproduction, each numerical gene is divided by two, the

PHOTO: KARL RABE

algorithm randomly "mutates" the results, and they're stored in virtual eggs and sperm. The program then pairs up the virtual gametes—with fitter robots providing a larger share—to concoct "genomes" for the next generation.

Often evolutionary roboticists don't build physical robots at all. Instead, computers simulate the robots and their behavior to determine which are fitter. Researchers quickly plow through thousands or millions of generations to optimize the design for a physical robot.

Long, however, likes to keep his evolving robots real, immersed in a physical environment. He grew up in Rochester, Michigan, but, as a descendent of New England whalers, felt the call of the sea at an early age. Hoping to become a marine biologist, he attended the tiny College of the Atlantic—enrollment 362—in Bar Harbor, Maine, where he worked for Sentiel "Butch" Rommel, a bioengineer who would take students out to slice up beached whales.

In graduate school at Duke University in Durham, North Carolina, Long studied fish, in particular the blue marlin, which can swim 80 kilometers per hour. He built a rig that would hold marlin backbones, begged from fish shops in Hawaii, and bend them back and forth. He expected that at higher bending frequencies, the backbones would become springier and absorb less energy. Instead he saw the opposite, suggesting the backbone works a bit like a shock absorber at high speeds.

Long was still pondering backbones when he arrived at Vassar in 1991. He wondered how the chainlike spine of vertebrates evolved from the sinewy notochord of invertebrates—an innovation that has happened at least three times in evolutionary history. Stiffer than a notochord, a backbone may have helped early fish swim faster and gather more food than their floppier peers, he hypothesized. Long had started to build mechanical models of fish, but his collaborator Kenneth Livingston, a cognitive scientist at Vassar, suggested developing them into autonomous robots.

To try to replicate backbone evolution, Long decided to mimic the tadpolelike larva of invertebrates of the genus *Botrylloides*, commonly called sea squirts. The larva resembles ancient invertebrate fish, such as the extinct eel-like *Haikouichthys*. Other researchers had traced the neural circuits that make sea squirt larvae spiral toward light, so the robots could be wired accurately and provide a reasonable model of the larva, which would stand in for the fish.

Thus the Tadro—short for tadpole robot—was hatched, in 2004. Its simplicity hit the sweet spot for research with undergraduates, who can't afford to spend years constructing a complex robot. Long's team built the first Tadros out of food containers, says Nicholas Livingston, chief engineer in Long's lab and Kenneth Livingston's son. "I remember feeling both clever and silly when I went to the local grocery store and bought a lot of Tupperware and plastic wrap," he says.

Thinking inside the box

Researchers hope to make their swimming robot evolve distinct control circuits (right) instead of one spaghettilike tangle (left).

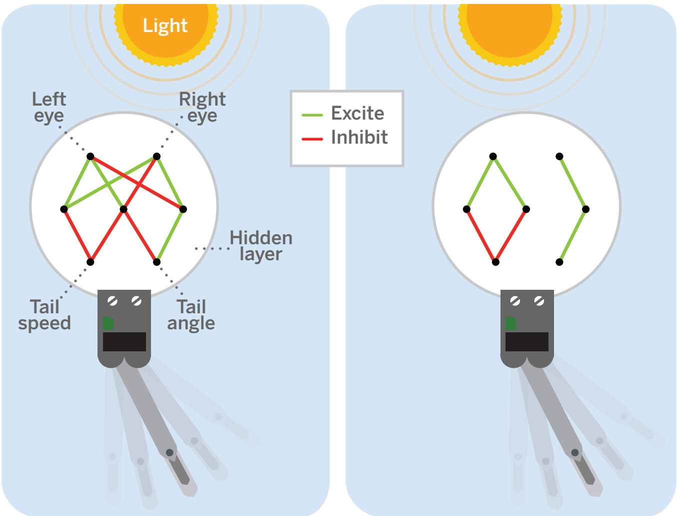

ILLUSTRATION: G. GRULLÓN/SCIENCE

COMPARED WITH COMPUTER MODELS or biological studies, robots have advantages for reprising evolution. Unlike a computer simulation, a real robot cannot leave out some subtle interaction with its environment or break the laws of physics. And robot experiments can be much faster than biological experiments, at least those with larger animals, says Jodi Schwarz, a bioinformaticist and collaborator at Vassar. "Even with mice, each generation is months," she says, "so it would take you forever."

Even so, studying evolution with robots can be tricky. In his first experiment, Long varied the length and stiffness of the Tadro's tail, molding new tails for each generation out of a tunable polymer. Researchers ran trials with three Tadros programmed to find and circle a light and assessed their fitness. Because the robots did not actually live or die—the gold standard for measuring fitness—the researchers had to find some other measure of fitness. They settled on a metric based on how fast and straight a Tadro swam, as assessed in a frame-by-frame video analysis.

But over 10 generations they observed no clear trend toward longer, stiffer tails. Instead, the tails' characteristics varied almost randomly. Why? Long and company had tripped on a pitfall of evolutionary robotics. They had defined their fitness metric to reward a Tadro for speed and penalize it for wobble. But in reality, faster Tadros wobbled more than slower ones—so, paradoxically, the fitness measure both rewarded and penalized a Tadro for being fast.

Even after correcting for that error and reanalyzing their data, they found that in some generations, selection favored bendy, slower Tadros. So Long's hypothesis was wrong: The race for food alone probably did not create a need for speed and account for the evolution of vertebrae, as he explains in his 2012 book, *Darwin's Devices: What Evolving Robots Can Teach Us About the History of Life and the Future of Technology*.

That negative result only spurred Long on. By 2007, he had revised his hypothesis. Perhaps vertebrae evolved not only to enable a fish to gather more food, but also to help it dash away from predators, he thought. To test that idea, Long and colleagues deployed a new version of Tadro, this time modeled after the extinct jawless fish *Drepanaspis gemuendenensis*, a vertebrate that also appeared to have had to fend off predators, as it had a hard shell. And this time researchers made two kinds of Tadros: predators and prey. The prey would still seek the light. However, they would also have infrared sensors that would trigger them to flee whenever a predator got too close.

The results of the new experiment were more in keeping with expectations. To keep the experiment manageable, tail length was fixed, but prey Tadros' tails could have different numbers of vertebralike rings. More vertebrae meant more stiffness, which presumably meant more speed. The researchers ran two trials of five and 11 generations each. And the Tadros did indeed evolve to have more vertebrae, from a starting average of 4.5 to an average of 5.5 or more—bolstering Long's hypothesis.

Both experiments were limited to a handful of generations and couldn't replicate the dramatic effects in the fossil record. Nevertheless, the results were clear, says Eiben, the computer scientist from VU University Amsterdam. "I'm amazed that they found such developments in just this few generations," he says. He credits Long's group with not just observing a trend, but also digging deeper and "asking themselves why."

NOW, LONG AND COLLABORATORS aim to study the evolution of the Tadro's "brain," hoping to induce the appearance of the distinct neural network circuits that they believe may aid further evolution. This time, the experiment also promises a practical payoff, says Josh Bongard, a computer scientist from the University of Vermont in Burlington who is collaborating with Long.

Bongard uses simulations to develop control circuitry for robots. But he has been frustrated to find that the process runs out of steam: The circuits become so interconnected that changing one connection requires rewiring the whole thing. Nature avoids that tangle by evolving modular circuits, and he hopes to learn how do to that, too. "The roadblock isn't that we don't have big enough computers," he says. "The roadblock is intellectual—we're not simulating evolution in the right way."

The new Tadro's neural network is rudimentary (see figure, p. 193). Inscribed on a reprogrammable chip, it has two "neurons" that take input from the eyelike light sensors on either side of the head. Two output neurons control, independently, the tail's angle and flapping rate. Inputs connect to outputs through a "hidden layer" of neurons. The map of that circuitry will be the heritable trait, and researchers hope Tadros will evolve distinct circuits for controlling tail speed and angle.

Learning from previous work, Long's team will simplify the experiment. Each Tadro will ply the water alone, and fitness will be determined by how much light it collects as measured by the light sensor in the middle of the head. The researchers may eventually allow some aspect of the Tadro's body, such as its tail, to evolve, too, as modularity may arise from the interplay of body function and the environment, Bongard says.

But first, the team must get the new Tadros running—a challenge for the four undergrads who will work for 2 years on the project. They had planned to make the Tadros entirely with a 3D printer over the summer. But the printed bodies leaked and sank, even after researchers tried to seal them with paint. "We had five different bodies painted with five different spray paints, and they all leaked," says Jessica Ng, a biology major at Vassar. "So we just gave up." The team is now using bodies machined from clear acrylic.

Long's students also face the daunting task of taking the data during the school year. Long, a passionate teacher who arrives at work at 7 a.m. and likes to squeeze in instruction whenever he can—he regularly reads aloud to his wife and two teenaged daughters at dinner—says he strives to help students succeed on their own terms. "My line to them is 'You're no good to us if you fail,'" he says. Instead of committing a number of hours per week to the project, the students will aim for monthly research milestones. If that's a recipe for periodic all-nighters, so be it, says John Loree, a physics major at Vassar: "Hey, it's college!"

It's tantalizing to imagine that Long's modest robots will evolve wildly. If they do, they might even provide insight into deep philosophical issues about the genome, says Kenneth Livingston, the cognitive scientist. "The genome is a set of instructions to do something, but it's not pure because it's about a particular world," he says. So, he says, the connection between the environment and the genome is "the beginning of meaning."

But it's also imaginable that things won't go so swimmingly. For example, the Tadros' evolution will depend in part on randomly tweaking connections within their neural networks. But most randomly wired networks may not work at all, leaving the Tadro dead in the water and evolution at a standstill—although simulations suggest that won't be a showstopper, Bongard says.

"We really have no idea what will happen," says Schwarz, the Vassar bioinformaticist. But that's part of the attraction of studying evolution with robots, says Pennock, the philosopher of science at Michigan State. "The thing that is underappreciated in this approach is that it's truly experimental," he says. "You're often surprised." ■

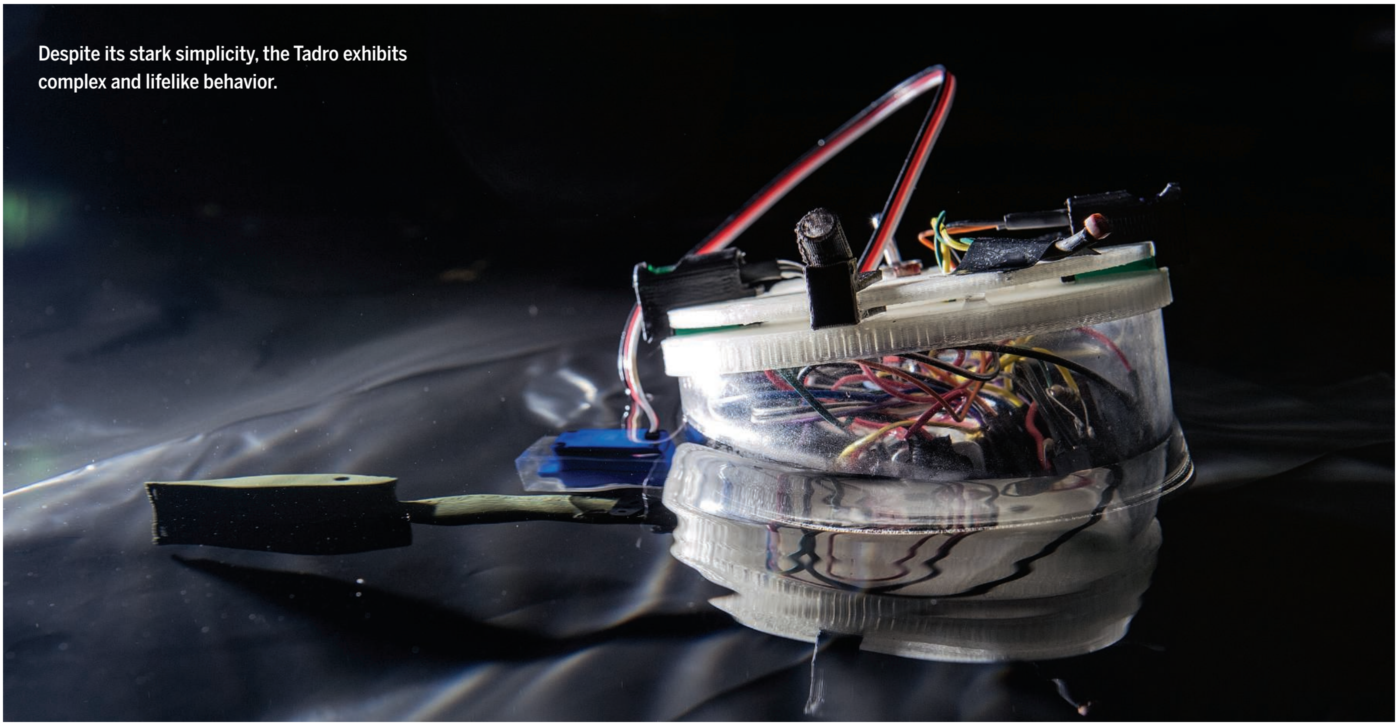

Despite its stark simplicity, the Tadro exhibits complex and lifelike behavior.

PHOTO: KARL RABE