Advanced mathematical techniques enable AUVs to survey large, complex and cluttered seascapes

November 8 2011, by Nancy Stauffer

Photo courtesy of MIT Sea Grant College Program

Since the 1970s, when early autonomous underwater vehicles (AUVs) were developed at MIT, Institute scientists have tackled various barriers to robots that can travel autonomously in the deep ocean. This four-part series examines current MIT efforts to refine AUVs' artificial intelligence, navigation, stability and tenacity.

Imagine dropping an underwater vehicle into the ocean and having it survey the ocean floor for debris from an accident or examine a ship's hull for signs of damage. Without any outside guidance or prior knowledge, the vehicle would traverse the target area in a methodical

fashion, never repeating itself or going astray, all the while generating a map that shows the surface of interest.

An MIT team has developed advanced mathematical techniques that enable such a scenario to occur — even when the area being examined is large, complex and cluttered, and the information coming from the vehicle's sensors is not always clear and accurate.

"A big problem for an autonomous underwater vehicle is knowing where it's been, where it is now and where it should go next — without any outside help," says John J. Leonard, a professor of mechanical and ocean engineering and a member of the MIT Computer Science and Artificial Intelligence Laboratory. Navigating underwater is tricky. Radio waves don't propagate through seawater, so an AUV can't use GPS as a guide. Optical methods don't work well. Computer vision is difficult, even for terrestrial robots; water reflects and refracts light in complex ways, and visibility may be poor due to murkiness and turbidity.

What's left? Sound waves, which can be monitored by acoustic sensors. To help an underwater vehicle navigate, a deepwater energy company may drop a network of acoustic transponders onto the seafloor. The vehicle exchanges acoustic "pings" with the transponders, generating data with which it can calculate its position. But sometimes the signal bounces off extraneous objects, producing inaccurate data. Sometimes several robots share multiple transponders, leading to confusion. And sometimes deploying enough transponders to cover a sufficiently large area is prohibitively expensive.

"So here's the challenge. You want to place the AUV at an unknown location in an unknown environment and, using only data from its acoustic sensors, let it incrementally build a map while at the same time determining its location on the map," Leonard says. Robot designers have studied the so-called mapping problem for decades, but it's still not

solved. As Leonard notes, it's a chicken-and-egg problem: You need to know where you are to build the map, but you need the map to know where you are.

To illustrate how robotic mapping works — and doesn't work — Leonard considers the aftermath of a hypothetical accident. The seabed is covered with debris, and officials need to figure out where it all is. Ideally they'd send down an AUV and have it cruise back and forth in a lawnmower-type pattern, recording information about where it is and what it sees.

One conventional way of accomplishing that task is using dead reckoning. The AUV starts out at a given position and simply keeps track of how fast and in what direction it's going. Based on that information, it should know where it is located at any point in time. But the calculations to determine its position quickly become wrong, and over time, the error grows "without bounds." Leonard likens it to mowing the lawn blindfolded. "If you just use dead reckoning, you're going to get lost," he says. Using expensive accelerometers, gyroscopes and other equipment will make the error grow more slowly, but not eliminate it entirely.

So how can an AUV use poor data from relatively inexpensive sensors to build a map? To tackle that problem, Leonard and his team have been using a technique called Simultaneous Localization and Mapping, or SLAM. With this approach, the AUV records information, builds a map and concurrently uses that map to navigate. To do so, it keeps track of objects it observes — in the accident example, say, a particular piece of debris on the seafloor. When the AUV detects the same object a second time — perhaps from a different vantage point — that new information creates a "constraint" on the current map. The computer program generating the map now adds that object and at the same time optimizes the map to make its layout consistent with this new constraint. The map

adjusts, becoming more accurate.

"So you can use that information to take out the error, or at least some of the error, that has accrued between the first time you saw that object and the next time you saw it," Leonard says. Over time, the program continues to optimize the map, finding the version that best fits the growing set of observations of the vehicle's environment.

In some cases, the AUV may see the same object again just a few minutes later. Identifying it as the same object is easy. But sometimes — especially when surveying a large area — the AUV may see the same object early on and then again much later, possibly even at the end of its travels. The result is a "loop closing" constraint. "That's a very powerful constraint because it lets us dramatically reduce the error," Leonard says. "That helps us get the best estimate of the trajectory of the vehicle and the structure of the map."

While SLAM has been in use for several decades, the Leonard group has made significant advances. For example, they've come up with new computational algorithms that can calculate the most likely map given a set of observations — and can do it at high speed and with unprecedented accuracy, even as new sensor information continues to arrive. Another algorithm can help determine whether a feature that the robot sees now is in fact the same one it saw in the past. Thus, even with ambiguous data, the algorithm can reject incorrect "feature matching" that would have made the map less rather than more accurate.

Finally, their methods ensure that uncertainty is explicitly addressed. Leonard emphasizes that SLAM may not produce a perfect map. "It's easy for a vehicle to get fooled by errors in the acoustic information," he says. "So we don't want to be overconfident. There's a certain inherent uncertainty to the sensor data, and it's important to get that uncertainty right. So we're not only building the map but also including the right

error bounds on it."

A problem of particular interest to Leonard is using AUVs to enable rapid response to accidents and other unforeseen events. For example, one challenge during the April 2010 Deepwater Horizon oil spill was determining whether there was a spreading plume of oil and if so, tracking where it was going. A network of AUVs working together could play a critical role in carrying out such tasks.

To that end, Leonard and his team are developing techniques that will enable AUVs to communicate with one another so they can navigate and collect information cooperatively. "If they can share information, they can accumulate data far more quickly than if they work alone," he says. "Together, they'll be able to sweep a large area and quickly produce the best possible map so that people can understand what's going on and develop and implement an effective response."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/*), a popular site that covers news about MIT research, innovation and teaching.*

 More information: *Next: Biomimetic pressure sensors help guide oceangoing vessels.*

Provided by Massachusetts Institute of Technology