

Microsoft researchers test AI-controlled soaring machine

August 17 2017, by Allison Linn



From left, Debadeepta Dey, Andrey Kolobov, Rick Rogahn, Ashish Kapoor and Jim Piavis prepare to launch a type of glider known as a sailplane in the desert in Hawthorne, Nevada. Credit: Microsoft

In the searing midday heat of the Nevada desert, a white Jeep Wrangler

heads down a desolate strip of dirt road, surrounded on either side by miles of sagebrush and sand.

As the Jeep bumps along, two members of a Microsoft research team, Jim Piavis and Rick Rogahn, steady themselves against the roll bar, their feet planted on the seats and their upper bodies jutting out of the open roof. They are scanning the bright blue sky, tracking a type of glider known as a sailplane.

The delicate black, white and red sailplane swerves wildly and unevenly at first, and then gradually begins to make wide, soaring circles.

A hawk appears next to it, following the same circular pattern.

"We've got a friend up there with us," says Piavis, head of mission readiness. "That's a good sign."

The 16 ½ -foot, 12 ½- pound sailplane has found a thermal, or an invisible column of air that rises due to heat. Soon, it is soaring through the sky, the Jeep reaching speeds up to 30 miles per hour as it flies down the dusty dirt road in hot pursuit of what the team is calling the infinite soaring machine.

Microsoft researchers have created a system that uses artificial intelligence to keep the sailplane in the air without using a motor, by autonomously finding and catching rides on naturally occurring thermals, similar to how many birds stay aloft.

"Birds do this seamlessly, and all they're doing is harnessing nature. And they do it with a peanut-sized brain," says Ashish Kapoor, a principal researcher at Microsoft.

The birds do it naturally. For a machine to do it requires a complex set

of AI algorithms that can identify things like air temperature, wind direction and areas where it is not allowed to fly. Then, the system must use other AI methods to take that information and make real-time predictions about where it might find its next ride on a thermal.

Taken together, that's much more complex than most AI systems people are using today for individual tasks like recognizing a face in a photo or the words in a conversation.

Kapoor says it's probably one of the few AI systems operating in the real world that's not only making predictions but also taking action based on those predictions.

It's still a work in progress, but Kapoor says the infinite soaring machine could eventually be used for all sorts of practical tasks, such as monitoring crops in rural areas or providing mobile Internet service in a place where there's no easy way to get needed connectivity.

"These can be your cellular towers someday," Kapoor says. "You don't need any ground infrastructure."

Eventually, the team says, the sailplane could even use solar or wind power to gather energy, theoretically making it possible for it to stay aloft indefinitely.

AI in the real world

The autonomous sailplane is useful in and of itself. But Andrey Kolobov, the Microsoft researcher in charge of the project's research and engineering efforts, says they also expect the work to apply to plenty of other increasingly sophisticated systems that rely on AI and will operate in real, unpredictable environments.

"For us, the sailplane is a testbed for technologies at the core of anything that will be considered intelligent in the next 10 years," he says.

For people to depend on AI to help them with things like driving their cars, keeping their homes secure or managing their busy daily schedules, these systems will need to reliably make complex decisions on the spot, based on variables such as traffic, noise, weather, other objects and even human emotions. What's more, they won't be able to make many mistakes – that would be costly and potentially dangerous.

"AI in the real world will have very little room for error, like our sailplane," Kolobov says.

In research terms, this kind of ability is called sequential decision making under uncertainty.

"It's really the question of, 'How do you plan for the future, several steps ahead?'" Kapoor says. "Computationally, that's a very hard problem."

AI theory in action

The sailplane they are testing in Nevada relies on a battery to run onboard computational equipment and controls such as the rudder, plus radios to communicate with the ground. It also has a motor so that a pilot can take over manual operation when necessary. But once it's up in the air, it's being designed to operate on its own, finding and using thermals to travel without the aid of the motor or a person.

"What we're trying to do is make sure the sailplane is fully autonomous and smart enough to change its course of action," Kolobov says.

To design the system, the team of researchers began with a framework for thinking about the problem called the partially observable Markov

decision process.

Kolobov, who co-authored a book on Markov decision process, says that's a model for making planning decisions in an environment in which you can't know everything. With the sailplane, the team combined that model with another AI approach, called Bayesian reinforcement learning, to create a way for the system to learn what it needs to know about its environment as quickly as possible, in order to make the right decisions.



Microsoft's aerial robotics research team prepares to launch a sailplane during a test flight in Hawthorne, Nevada. Credit: Microsoft

The team also is using what's called Monte Carlo tree search, which is a way for AI to look for the most promising course of action.

The sailplane's AI system is divided into two parts: The high-level and the low-level planner.

The high-level planner takes all the factors of the environment into account and tries to create a policy for where the sailplane should go to look for thermals. It gets better at making those predictions as time goes on, based on the information the sailplane collects each time it goes up in the air.

"For the high-level planner, experience matters," Kolobov says. "The system will perform better on Friday than on Thursday because it incorporates information based on past flights."

The low-level planner is the part that is using Bayesian reinforcement learning to detect and latch onto thermals in real time, based on data from the sailplane's sensors. Think of that as learning by doing.

Reality testing

To build those AI algorithms required months of work in offices at Microsoft's Redmond, Washington, campus.

Once the weather turned warm and the thermals got better, the team conducted limited test flights on a farm near the company's headquarters. But it wasn't until they arrived at the tiny airfield in Hawthorne, Nevada, in mid-August that they really got a chance to see how all their theories would play out in the real world.

On the day that they spot the hawk, the team has been out on this stretch of dirt road for four days, battling dusty conditions and the blazing sun

as they launch their sailplanes again and again.

Behind the Jeep, an extra-long burgundy Ford Expedition has been set up as a makeshift office, crammed with multiple computers, tools for field repairs and four other members of the research team. The blasting AC does little to combat the heat from the desert and the large amounts of computer equipment.

As the sailplane flies above them, research interns Iain Guilliard and Sangwoo Moon – both of whom wrote core algorithms that help power the system – use four laptops to track its progress and monitor how it is taking advantage of conditions including thermals.

Every few seconds, Guilliard calls out the sailplane's flight parameters, while those in the Jeep scan the air for physical proof of the sailplane's position. Kolobov and Debadepta Dey, a Microsoft researcher who built the sailplane's thermal prediction module, take turns driving the movable office after the Jeep and the sailplane.

In addition to catching thermals, the system needs to plan its actions to avoid certain obstacles, such as nearby mountains, a large lake and scores of munitions that the U.S. Army stores in the area near the test flight site.



Rogahn monitors the sailplane's flight path during a test flight near Hawthorne, Nevada. Rogahn is the team's human pilot, meaning he is responsible for manually launching and landing the aircraft, and for intervening if it is necessary for safety. Credit: Microsoft

The messy real-world environment, with all the quirks and obstacles that would be impossible to predict in the office, is providing the perfect testing ground for the system.

Many of the problems they encounter, and have to solve, have nothing to do with AI at all. A glitch in the communications system creates problems with the largest sailplane they are testing, and the team suspects that magnetic particles in the dust on the road are responsible for causing damage to an electronic component.

During one test, when they try to launch a sailplane using a bungee cord attached to the Jeep after a motor failure, the bungee cord breaks.

With each setback, the team patiently sets to work to find a solution: Fixing the motor, repairing the bungee cord, finding a new battery, pulling out another sailplane when one is taken out of commission.

Kolobov says these real-world problems are exactly what drew many people to this project. It's also why they aren't too discouraged by each unexpected thing they encounter.

"This is why reality is different from simulation," Kolobov says. "And that's what we came here for. We came here to learn, and it's not necessarily the things we expected to learn."

'The algorithm is doing better than me'

It's late on a Friday afternoon when the team gathers at the tiny Hawthorne airport for one last debrief.

About an hour earlier, on the team's last flight of the trip, the algorithms had performed exactly how they'd anticipated, sending the sailplane soaring through the air – before the battery unexpectedly died, the sailplane nosediving toward the sand and sagebrush with alarming speed.

Rogahn, whose job was to be the human pilot – providing a backup to the AI and maneuvering it to safety when necessary – was able to regain control at the very last moment, preventing a crash by perhaps inches.

"It's not polite to wear a cowboy hat indoors, but I feel like I've ridden a couple of bulls," Rogahn says at the debrief, drawing laughter from the group.

But then Rogahn grows serious, removing the hat and acknowledging how far they've come during the long days in the desert.

"We've reached the point this summer where it's actually doing better than me," he says. "The algorithm is doing better than me as a sailplane pilot."

Provided by Microsoft

Citation: Microsoft researchers test AI-controlled soaring machine (2017, August 17) retrieved 20 March 2024 from

<https://phys.org/news/2017-08-microsoft-ai-controlled-soaring-machine.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--