Sikorsky has successfully integrated its Matrix Technology into its fly-by-wire SARA S-76B, its optionally piloted UH-60A, and its ALIAS autonomy kit. While Matrix incorporates nondeterministic algorithms in its outer loop for non-flight-critical functions, such methods will likely play an even greater role in future autonomous aircraft. (Sikorsky photo)

Intelligent Black Boxes

Advancements in artificial intelligence hold great promise for autonomous aircraft — but how do you certify a system you can't fully explain?

By Elan Head

For any primary helicopter flight instructor, a student's first controlled liftoff to a hover is a milestone event. So it was for Sikorsky experimental test pilot Mark Ward, with one important difference — his "student" was the helicopter itself.

As the chief test pilot for the Sikorsky Autonomy Research Aircraft (SARA), an S-76B retrofitted with fly-by-wire flight controls, Ward has been along for the ride as SARA has progressed from straight-and-level flight, to takeoffs and landings, to more complex maneuvers, such as an approach to a moving barge.

"Picture a big steel barge being towed by a towboat at about 8–10 kt [15–20 km/h], off of Groton, Connecticut. We had it painted up like a helideck … and we did an autonomous takeoff, we intercepted a course to that ship and we did an approach," Ward recalled.

"I've done flights where my hands never once touched the controls. It's pretty remarkable when you think about it. And there's nobody on the ground flying this thing either. It truly is a self-contained system flying on its own."

Functionally, aircraft like SARA are the logical progression of the trend towards increasing automation in the cockpit. Behind the curtain, however, they represent a huge technological leap.

The programming that tells today's autopilots to track to a specific waypoint or descend on a glide slope is essentially deterministic: if this, then that. But aircraft operate in complex stochastic environments. If you want to take autopilots to the next level — giving them the ability to perceive and avoid obstacles, or identify and evaluate landing zones — it is virtually impossible to write deterministic algorithms that cover every contingency.

Recent advancements in artificial intelligence (AI) have allowed autonomous vehicle developers to make tremendous strides in tackling these kinds of problems, and could lead to even greater gains in the future. But these nondeterministic methods come with a caveat: they aren't entirely explainable or predictable.

That poses a particular challenge in the aviation industry, in which safety is paramount and demonstrating reliability is an essential aspect of certification. As aircraft manufacturers experiment with ever-greater levels of autonomy, they must balance the promise of AI against considerations of certifiability and public acceptance.

New Tools

First, some notes on terminology. The comprehensive form of artificial intelligence portrayed in science fiction — from *2001: A*

"...nondeterministic methods come with a caveat: they aren't entirely explainable or predictable."

Space Odyssey to *Westworld* — is known as "general AI," and most experts believe it's still many years in the future. In this article, "AI" will refer to "narrow AI," computerized systems that can accomplish specific tasks that call for reason or decision-making.

AI has exploded as a research field in the past decade. As a 2016 report from the US Office of Science and Technology Policy (OSTP) National Science and Technology Council (NSTC) — entitled "Preparing for the Future of Artificial Intelligence" — explained, this has been due to the convergence of three factors: 1) new sources of big data, which have provided material for 2) improved machine learning approaches and algorithms, which have been enabled by 3) more powerful computers. AI is already in widespread use in applications including voice and image recognition, and email spam filtering.

In machine learning, a computer algorithm is presented with data and proceeds to identify patterns that can explain the data or predict future data. A subfield of machine learning is "deep learning," which uses layered structures called artificial neural networks.

This approach of learning from data is a departure from the more established "expert systems" approach to AI, in which programmers rely on insight from human experts to codify rules for decision-making. As Airbus Helicopters head of engineering Jean-Brice Dumont noted at Forum 73 in May, contemporary autopilots "are built more like expert systems," in that they "are not more than the applications of precise and predetermined rules."

And the evolution of AI away from rule-based programming has yielded new approaches to the innumerable challenges associated with navigating a vehicle through an unpredictable environment.

An unskilled Marine operator directs an unmanned Boeing H-6U — equipped with the initial Autonomous Aerial Cargo/Utility System (AACUS) capability — from a tablet in April 2014. (Aurora photo)

"As you may imagine, there's a whole bunch of problems, especially in perception, that don't necessarily lend themselves to a deterministic solution," said Sikorsky director of autonomous programs Igor Cherepinsky, who has led development of SARA. "The further away you go from basic vehicle control, the less deterministic it gets, and the higher levels of behavior drive those nondeterministic methods and applications."

As an example, convolutional neural networks (CNNs) have revolutionized the computational pattern recognition process, according to Jae-Woo Choi, autonomy and perception program lead at Aurora Flight Sciences. Previously, researchers working in perception had to laboriously engineer methods for extracting and classifying features in visual scenes. Today, he said, "features in scenes are trained by a CNN using training examples; this technique is especially powerful because the convolution kernels scan an entire image and few parameters need to be learned."

James Paduano, Aurora's director of autonomy, added, "For tasking, route and trajectory planning, it's very hard to beat the nondeterministic methods." Although Aurora is not currently looking at the newest learning/AI methods for these problems, he said "by virtue of our methods for incorporating nondeterministic algorithms at this level, we would be open to the possibility."

Meanwhile, deep learning techniques will likely become increasingly important for understanding the wealth of data

Aurora Flight Sciences is using nondeterministic methods — including the algorithms RRT* and RRT# — in its work on AACUS. Aurora has demonstrated its modular Tactical Autonomous Aerial Logistics System (TALOS) on a Bell UH-1H and other helicopters. (Aurora photo)

generated by autonomous platforms, noted Michael McNair, an engineer with Bell Helicopter's Innovation team. "With autonomous systems and their onboard sensor packages, a tremendous amount of data gets generated, and it's important to be able to go through that data and analyze that data," he said. "Basically you're looking for patterns, you're looking for trends, and that's where deep learning would come into play."

Managing Uncertainty

While models generated through machine learning can be remarkably effective at interpreting data, they can also be dauntingly complex. Consequently, as the 2016 NSTC report observed, "it is typically not possible to extract or generate a straightforward explanation for why a particular trained model is effective."

That poses a real problem for aircraft certification, in which stringent safety requirements don't permit much ambiguity.

Sikorsky is looking at a variety of ways to incorporate autonomy into its current manned platforms. For example, in the event of an engine failure, a semi-autonomous helicopter could identify a suitable landing zone faster than a human pilot. Such image recognition capabilities leverage nondeterministic processes. SARA autonomously landed in a clearing in 2013. (Sikorsky photo)

"At the top level, what all the certifying agencies want to know, and want you to prove, is that if you say the system is going to do X, it in fact does X," said Cherepinsky. "So obviously, as you give the system more and more autonomy — pun intended — the position of the X starts to grow, and being able to prove beyond the shadow of a doubt that it's going to do that gets harder and harder."

For Sikorsky, which aims to introduce autonomous capabilities into its current civilian and military product line, certification concerns have shaped its approach to SARA and the aircraft's underlying MATRIX Technology, a capability toolkit that includes hardware and software as well as multi-spectral sensors that enable scalable automation.

"The challenge is obviously how do you certify something that doesn't have deterministic behavior," said Chris Van Buiten, Sikorsky's vice president of Innovations. "And the first-order answer to that is to make sure that that is not a flight-critical function. So imagine having advisory information be nondeterministic, and flight-critical being more deterministic."

That approach is evident in the Defense Advanced Research Projects Agency (DARPA) Aircrew Labor In-Cockpit Automation System (ALIAS) program, for which Sikorsky was recently awarded a Phase 3 contract. ALIAS envisions a drop-in autonomy kit for existing manned aircraft — including the UH-60 Black Hawk — that would enable them to be flown with reduced onboard crew.

Dr. Graham Drozeski, a program manager in DARPA's Tactical Technology Office, is overseeing the ALIAS program. "I've worked with helicopter autonomy programs for 15 or 20 years, and I've always seen very promising results in terms of the capabilities of unmanned helicopters; but on none of the programs I've worked on previously have we succeeded in transitioning those technologies in a certified manner to man-rated helicopters," he said.

"This is the first program that I think has an opportunity to do that. And we've taken steps and Sikorsky has taken steps along the way — when they chose their algorithms, when they chose their hardware, when they chose which aspects of the aircraft operation they're going to conduct autonomously — to ensure that there is a path to certification on a large helicopter."

While the current certification framework may be able to accommodate limited applications of nondeterministic methods, more extensive adoption of these techniques will likely require new verification and validation methodologies.

"Fundamentally, any kind of certification requires being able to check behaviors and processing against a certain set of criteria. And when it comes to aircraft certification, your criteria is formed by regulations and standards," said McNair. "There are different standards bodies that are working on different aspects of this whole problem area of autonomous vehicles. Their work is incomplete; it is definitely still in progress. So without the standards work being very mature, it's going to be a little bit before we actually see mature certification requirements."

An Intelligent Future

It's worth keeping in mind that today's aircraft already incorporate nondeterministic systems in the form of human pilots. Pilots undergo extensive training, but how they actually react in a given situation could depend on a host of unpredictable factors — from how they slept the night before to what they ate for breakfast. And because pilots may not be fully aware of these factors themselves, their post hoc explanations for why they made certain decisions might be incomplete.

With the ALIAS program, DARPA aims to leverage autonomy to reduce workload for human pilots, allowing them to "use their cognitive capability to do much more than manage the systems on the aircraft and certainly more than moving sticks," said Graham Drozeski. The Aurora ALIAS kit is shown here in a Boeing 737 simulator; the Sikorsky ALIAS did not use hardware manipulators, but rather tied in directly to the fly-by-wire controls. (Aurora photo)

Sikorsky's Matrix may be used on future civil and military aircraft, as demonstrated on the Optionally Piloted Black Hawk in 2014. (Sikorsky photo)

"Today, when you fly on an airliner, there's two pilots up there, and I dare anybody to explain their decisions per se," said Cherepinsky. "But they've gone through a training regimen, and they've gotten examined, to whatever degree we can examine them, and we think they're qualified. So now whether a similar thing will ever be acceptable with any machine, is a much more cultural question really than anything else."

McNair noted that public acceptance isn't simply about marketability — it can also shape how certification is applied through legislation. "It's important to realize that certification is more than just standards and regulation; it includes other aspects that are certainly influencing factors," he said.

How willing operators and passengers will be to embrace AI in the cockpit will depend, in part, on how the technology evolves. Cherepinsky, for one, believes that "anyone who thinks that current machine learning, deep learning, is the end of it all is probably not paying attention to history."

DARPA is currently leading several initiatives to advance AI at a fundamental level. One of them is the Explainable AI (XAI) program, which aims to create a suite of machine learning techniques that produce more explainable models, enabling human users "to understand, appropriately trust, and effectively manage the emerging generation of artificial intelligence partners," according to DARPA's website.

Another is the Lifelong Learning Machines (L2M) program, which will focus on developing next-generation machine learning technologies that can continuously apply the results of past experience to new data or situations. The program will also study how biological systems learn and adapt, and whether those principles and techniques can be applied to machine learning systems.

"When it comes to the sponsors of this kind of exploration, of course you have government interest, whether it's through DARPA or any of the other [US military] services, NASA, and other agencies throughout the US government … And of course there are a number of private companies that have put a tremendous amount of money and resources into machine learning," said McNair. "It might be for their particular application, but in general it does further the technology as a whole."

Van Buiten suggested that the progress of AI in autonomous aircraft development may follow a "crawl-walk-run" model as the technology develops and certification standards and public acceptance evolve in turn.

"This first level, it's going to be largely deterministic with maybe some nondeterministic outer loop functions. And then the 'walk' generation is going to be more nondeterminism with improved verification and validation tools," he said.

As for the final "run," a fully AI-enabled, nondeterministic system would make possible military and commercial capabilities that today we can only imagine. But, said Van Buiten, "who knows how long it could take to get there."

About the Author

Elan Head is a freelance writer and serves as the special projects editor for *Vertical Magazine*, having previously served as its editor-in-chief. She is also a commercial helicopter pilot and flight instructor who has flown in the US, Canada and Australia. She was the moderator of the AHS Forum 73 "Straight Talk from the Top" CEO Panel.