The Rise of A.I. Fighter Pilots

Artificial intelligence is being taught to fly warplanes. Can the technology be trusted?

By Sue Halpern   January 17, 2022

Implementing new technology will mean convincing humans to cede control.Illustration by Karolis Strautniekas

On a cloudless morning last May, a pilot took off from the Niagara Falls International Airport, heading for restricted military airspace over Lake Ontario. The plane, which bore the insignia of the United States Air Force, was a repurposed Czechoslovak jet, an L-39 Albatros, purchased by a private defense contractor. The bay in front of the cockpit was filled with sensors and computer processors that recorded the aircraft's performance. For two hours, the pilot flew counterclockwise around the lake. Engineers on the ground, under contract with *DARPA*, the Defense Department's research agency, had choreographed every turn, every pitch and roll, in an attempt to do something unprecedented: design a plane that can fly and engage in aerial combat—dogfighting—without a human pilot operating it.

The exercise was an early step in the agency's Air Combat Evolution program, known as *ACE*, one of more than six hundred Department of Defense projects that are incorporating artificial intelligence into war-fighting. This year, the Pentagon plans to spend close to a billion dollars on A.I.-related technology. The Navy is building unmanned vessels that can stay at sea for months; the Army is developing a fleet of robotic combat vehicles. Artificial intelligence is being designed to improve supply logistics, intelligence gathering, and a category of wearable technology, sensors, and auxiliary robots that the military calls the Internet of Battlefield Things.

Algorithms are already good at flying planes. The first autopilot system, which involved connecting a gyroscope to the wings and tail of a plane, débuted in 1914, about a decade after the Wright brothers took flight. And a number of current military technologies, such as underwater mine detectors and laser-guided bombs, are autonomous once they are launched by humans. But few aspects of warfare are as complex as aerial combat. Paul Schifferle, the vice-president of flight research at Calspan, the company that's modifying the L-39 for *DARPA*, said, "The dogfight is probably the most dynamic flight profile in aviation, period."

A fighter plane equipped with artificial intelligence could eventually execute tighter turns, take greater risks, and get off better shots than human pilots. But the objective of the *ACE* program is to transform a pilot's role, not to remove it entirely. As *DARPA* envisions it, the A.I. will fly the plane in partnership with the pilot, who will remain "in the loop," monitoring what the A.I. is doing and intervening when necessary. According to the agency's Strategic Technology Office, a fighter jet with autonomous features will allow pilots to become "battle managers," directing squads of unmanned aircraft "like a football coach who chooses team members and then positions them on the field to run plays."

Stacie Pettyjohn, the director of the Defense Program at the Center for a New American Security, told me that the *ACE* program is part of a wider effort to "decompose our forces" into smaller, less expensive units. In other words, fewer humans and more expendable machines. *DARPA* calls this "mosaic warfare." In the case of aerial combat, Pettyjohn said, "these much smaller autonomous aircraft can be combined in unexpected ways to overwhelm adversaries with the complexity of it. If any one of them gets shot down, it's not as big of a deal."

All told, the L-39 was taken up above Lake Ontario twenty times, each sortie giving the engineers and computer scientists the information they need to build a model of its flight dynamics under various conditions. Like self-driving cars, autonomous planes use sensors to identify discrepancies between the outside world and the information encoded in their maps. But a dogfighting algorithm will have to take into account both the environment and the aircraft. A plane flies differently at varying altitudes and angles, on hot days versus cold ones, or if it's carrying an extra fuel tank or missiles.

"Most of the time, a plane flies straight and level," Phil Chu, an electrical engineer who serves as a science adviser to the *ACE* program, explained. "But when it's dogfighting you have to figure out, O.K., if I'm in a thirty-degree bank angle, ascending at twenty degrees, how much do I have to pull the stick to get to a forty-degree bank angle, rising at ten degrees?" And, because flight is three-dimensional, speed matters even more. "If it's flying slowly and you move the stick one way, you get a certain amount of turn out of it. If it's flying really fast and you move the stick the same way, you'll get a very different response."

In 2024, if the *ACE* program goes according to plan, four A.I.-enabled L-39s will participate in a live dogfight in the skies above Lake Ontario. To achieve that goal, *DARPA* has enlisted three dozen academic research centers and private companies, each working on one of two problem areas: how to get the plane to fly and fight on its own, and how to get pilots to trust the A.I. enough to use it. Robert Work, who was the Deputy Secretary of Defense during the Obama Administration, and pushed the Pentagon to pursue next-generation technologies, told me, "If you don't have trust, the human will always be watching the A.I. and saying, 'Oh, I've got to take over.' "

There is no guarantee that *ACE* will succeed. *DARPA* projects are time-limited experiments, typically lasting between three and five years. Schifferle, at Calspan, told me, "We're at the 'walk' stage of a typical 'crawl, walk, run' technology maturation process." Still, it seems increasingly likely that young pilots will one day wonder how their fighter jet acquired the skills of a Chuck Yeager. When they do, they will be told about a refurbished Soviet-era warplane that was flown high above Lake Ontario by old-school pilots who were, in a way, writing their own obituaries.

As part of the effort to devise an algorithm that can dogfight, *DARPA* selected eight software-development companies to participate in the AlphaDogfight Trials, an A.I. competition that culminated with three days of public scrimmages in August, 2020. The prize was a flight helmet worn by Colonel Dan (Animal) Javorsek, who was in charge of the program until he returned to the Air Force last year. The contest was supposed to be held in front of a live audience near Nellis Air Force Base, in Nevada, but the pandemic relegated the action to an online event, hosted by the Applied Physics Lab at Johns Hopkins, and broadcast via a YouTube channel called *DARPA*tv. Justin (Glock) Mock, an F-16 pilot, offered play-by-play commentary. At one point, he told the five thousand or so viewers that the objective was simple: “Kill and survive.”

Each team took a slightly different approach with its A.I. agents, as the algorithms are called. EpiSci, a defense contractor based in Poway, California, mounted an effort led by Chris Gentile, a retired Air Force test pilot. The company broke the problem down into component parts, and used Gentile’s flight expertise to solve each step. “We start at the lowest level,” Gentile told me. “How do you control the airplane? How do you fly it and direct it to go left and right, all the way up to what tactics should we use?”

PhysicsAI, in Pacifica, California, fielded a four-man squad who knew next to nothing about aerial combat. They used a neural-network approach, enabling the system to learn the patterns of a successful dogfight and mathematically arrive at the maximum probability of a good outcome. “The problem we have to solve is like playing chess while playing basketball,” John Pierre, PhysicsAI’s principal investigator, said. “You’re taking shots while making split-second decisions, and it needs to be done in a way that human pilots can interpret what’s going on.”

During each contest, the A.I. agents, represented by dime-size airplane avatars, moved around a screen at a stately pace, mimicking the flight dynamics of an F-16. It wasn’t exactly “Top Gun,” but the algorithms were doing something that would have been impossible a year earlier: interacting with each other and adjusting their tactics in real time. As the agents battled it out, Mock compared the action to “a knife fight in a phone booth.”

“And this hall reminds us we have always been very rich and quite ugly.”

Cartoon by Maddie Dai

In the decisive scrimmage, on day three, Falco, an A.I. agent created by Heron Systems, a boutique software company based in Virginia, competed against an A.I. agent developed by Lockheed Martin, the country’s largest defense contractor. The matchup drew the obvious David and Goliath comparisons—though this David had gone through about the same number of computer iterations as a pilot who trained all day, every day, for thirty-one years. After a few tightly fought rounds, Heron’s Falco emerged victorious. But there was a final contest: a seasoned F-16 pilot was going to take on Falco.

The pilot, dressed in an olive-green flight suit, sat in a high-backed gaming chair, his face obscured by a virtual-reality headset. He was identified only by his call sign, Banger. (His identity was concealed for “operational security.”) He’d trained with the team at A.P.L. beforehand, learning how to use the controls to guide his plane, and the V.R. headset to track his opponent’s vector of attack.

On a split screen, viewers could see what Banger saw from the cockpit. Another screen displayed a visual representation of the fight, as the planes—yellow for Banger, green for Falco—jockeyed for the best angle. About a minute in, each team aggressively rolled its aircraft, and Banger evaded the A.I. by dropping down to ten thousand feet. Falco came around and got off a series of good shots. Banger was down to four lives.

In the end, Banger failed to survive a single skirmish. He said, “I think technology has proven over the past few years that it’s able to think faster than humans and react faster in a precise pristine environment.” Banger also suggested that artificial intelligence might execute tactical maneuvers that pilots had been trained to avoid, such as flying too close to enemy aircraft and moving at speeds that would tax a human body. "I may not be comfortable putting my aircraft in a position where I might run into something else," he said. "The A.I. would exploit that."

Mock seemed pleased with the outcome. "You could look at this and say, 'O.K., the A.I. got five, our human got zero,' " he told viewers. "From the fighter-pilot world, we trust what works, and what we saw was that in this limited area, this specific scenario, we've got A.I. that works." (A YouTube video of the trials has since garnered half a million views.)

Brett Darcey, who runs Heron, told me that the company has used Falco to fly drones, completing seventy-four flights with zero crashes. But it's still unclear how the technology will react to the infinite possibilities of real-world conditions. The human mind processes more slowly than a computer, but it has the cognitive flexibility to adapt to unimagined circumstances; artificial intelligence, so far, does not. Anna Skinner, a human-factors psychologist, and another science adviser to the *ACE* program, told me, "Humans are able to draw on their experience and take reasonable actions in the face of uncertainty. And, especially in a combat situation, uncertainty is always going to be present."

In early May, I visited the Operator Performance Lab, at the University of Iowa, where members of the *ACE* program had gathered for a demonstration. O.P.L. is the creation of Tom Schnell, a Swiss-born professor of industrial and systems engineering. In his spare time, Schnell flies loops and rolls in an aerobatic plane above the cornfields of Iowa, but his expertise was, initially, in ground transportation. In the late nineties, a luxury-car company—Schnell won't say which one—asked him to develop a way to measure how much people enjoyed driving its vehicles. Schnell attached sensors to drivers' faces to detect the movement of small muscles around the mouth and eyes, which might indicate smiling or frowning, and electrocardiogram leads to monitor their heart. "I told them that if I was going to do this work they'd have to send me a fun car," Schnell said of his early client. "And they did."

Schnell soon found that each sensor came with its own proprietary data-collection system, which made it nearly impossible to analyze all the information at once. He built a common framework, which he named the Cognitive Assessment Tool Set, and began collecting the physiological data of people who operated all kinds of machinery. "They could be train engineers, or helicopter pilots, or people driving cars," he said. The face sensors supplied one set of data points. So did a device that analyzed galvanic skin response—how much a subject was sweating. Another tool looked at blood-oxidation levels, which served as a proxy for mental workload.

In 2004, Schnell persuaded his department chair at the University of Iowa to buy O.P.L.'s first aircraft, a single-engine Beechcraft Bonanza. Within a few years, he had acquired a jet, and commercial airlines and the Air Force hired him to conduct studies on their pilots. "We did a lot of work on spatial disorientation," Schnell said. This involved things like having pilots close their eyes during aerial maneuvers and then try to fly straight once they'd opened them again. By the time *DARPA* put out its request for proposals for the *ACE* program, in 2019, Schnell's laboratory had more than a decade of experience capturing the physiological responses of pilots.

Persuading pilots to hand over the controls may prove even more elusive than developing A.I. that can dogfight. "It's probably the paramount challenge we're trying to tackle," Ryan Hefron, the current *ACE* program manager, told me. Hefron, a thirty-eight-year-old lieutenant colonel with a doctorate in computer science, came to *DARPA* in 2021 from the Air Force Test Pilot School, where he was an instructor. As an example, he mentioned "Auto-*GCAS*," an automated ground-collision-avoidance system that grabs the controls if a plane is in imminent danger of crashing. During testing, Auto-*GCAS* had a tendency to pull up suddenly without cause—what Hefron called "nuisance fly-ups." The system has since saved at least eleven lives, but test pilots remained wary of it for years because of these early setbacks.

"There's a saying in the military," Peter Hancock, a psychology professor at the University of Central Florida who studies the effect of trust on technology adoption, told me. "Trust is gained in teaspoons and lost in buckets." It's not just an issue in warfare. In the most recent surveys conducted by the American Automobile Association, about eighty per cent of respondents said that they were not comfortable with the idea of autonomous vehicles. "Most of the drivers say that they want the current systems to work better before they can trust a fully self-driving system," Greg Brannon, the director of automotive engineering at AAA, told me. "The percentage hasn't moved much despite a lot of advances in technology, and that's pretty shocking."

To assess trust, psychologists typically administer surveys. "No one has ever come up with an objective measure of trust before," Skinner said. *DARPA* hired SoarTech, an A.I. research-and-development firm based in Ann Arbor, to build a "trust model," which aims to verify self-reported trust with the hard data from O.P.L.'s Cognitive Assessment Tool Set. "I think that's how you do good science," Schnell told me. "You take the best building blocks you have and put them together to answer very difficult questions. *DARPA* actually stepped up to the plate and said we want to know: 'Are you trusting the avionics?' "

One of O.P.L.'s hangars, at the Iowa City Municipal Airport, was filled with secondhand aircraft that Schnell had purchased and retrofitted: two L-29 Delfins, a smaller cousin of the L-39, painted a glossy Hawkeye yellow; a hulking Soviet helicopter, purchased for about the cost of a Cadillac Escalade, upgraded with a full-color night-vision system that Schnell built himself. At the far end of the hangar was the simulated cockpit of a 737 jet which was the size of a studio apartment.

An Air National Guard pilot, on loan to O.P.L. for the day, lowered himself into another simulator, a rectangular metal shell that Schnell called "the bathtub." Schnell hooked him up to electrocardiogram leads, in order to gather some baseline data. Until that morning's briefing, the pilot knew only that he would be participating in a *DARPA* research project. Even now, as he adjusted his V.R. headset and fidgeted with controls that replicated an F-16's, all he'd been told was that artificial intelligence would be controlling the plane while he played a rudimentary video game broadcast on his display panel. (A separate *ACE* effort is developing a more complex version.)

The game simulated the battle-management tasks that pilots are expected to conduct in the future; to win, the pilot's eight blue planes had to shoot down eight red enemy planes. An eye tracker inside his helmet would measure when and for how long he looked up to see what the A.I. was doing, which could be considered an expression of distrust. He did not know that some of the simulated skirmishes were primed for him to win and others to put him and his aircraft in jeopardy. But, if he felt that the A.I. was about to do something dangerous, he had the option of stopping the engagement by "paddling off." This, too, would demonstrate a lack of trust.

Ultimately, the idea is to supply pilots with more information about the A.I.'s next move, in order to elicit the appropriate level of trust. Glenn Taylor, a senior scientist at SoarTech, told me, "We're building visual and other interfaces into the system to let the pilot know what the A.I. is doing, and give him or her enough information, with enough time, to know whether or not to trust it." The researchers called this relationship "calibrated" trust. Phil Chu, one of the *ACE* program's science advisers, told me, "If we can show pilots what the A.I. is going to do in the next four seconds, that's a very long time."

Trust will also be crucial because, with planes flying at speeds of up to five hundred miles an hour, algorithms won't always be able to keep pilots in the loop. Hancock, the U.C.F. professor, calls the discrepancy in reaction time “temporal dissonance.” As an analogy, he pointed to air bags, which deploy within milliseconds, below the threshold of human perception. “As soon as you put me in that loop,” he said, “you’ve defeated the whole purpose of the air bag, which is to inflate almost instantaneously.”

In the “bathtub” at O.P.L., a computer relayed what the pilot was seeing in his goggles. As he turned his head to the right, a wing came into view; when he looked down, he could see farmland. A radar screen at the front of the cockpit kept track of the adversary, which, in the first skirmish, quickly gained an advantage, coming at the pilot from behind and preparing to take a shot. “Paddle,” the pilot called out, ending the skirmish. The computer was reset. One of Schnell’s graduate students, who helped design the experiment, counted down from three, then called “Hack” to start the next contest.

Forty minutes later, as the pilot left the simulator, he was greeted by Katharine Woodruff, a researcher working with SoarTech. Woodruff asked him about an incident in which he stopped the encounter even though he was not in imminent danger. “I had two decisions I could make,” he said. “To let it ride and see what happens or paddle off.” After a moment, he added, “I assessed that the bandit was starting to turn towards me. And so I paddled off.”

For the most part, Woodruff told me, the pilots in the study trusted the A.I. when it behaved appropriately and took over when it didn’t. There were a few exceptions: a pilot who had recently ejected from his plane was deeply suspicious of the technology. The thirty-year-old pilot whom I had observed thought that the autonomy “was cool,” but he paddled off even when his plane had the potential to achieve a good offensive angle. “I wanted to basically figure out my limits with the A.I.,” he told Woodruff. “What is too conservative, and what is going to get me killed. And then find that happy medium.”

Schnell’s graduate student, who can’t be named because he’s on active duty in the military, came over to listen to the debriefing. “You would be the perfect example of someone we’d need to influence, because—and I do not mean this to be rude at all—you completely violated the construct of the experiment,” he told the pilot. “You were deciding to not let the A.I. do the job that it’s put there to do, even though it was actually performing fine in the sense of not getting you killed. If we want to make you a battle manager in thirty years, we’d need to be able to push that behavior in the opposite direction.”

In 2017, the Future of Life Institute, an advocacy group focussed on “keeping artificial intelligence beneficial,” which counts Elon Musk as a member of its advisory board, released “Slaughterbots.” The short film imagines a world in which weaponized quadcopters about the size of a smartphone target political dissidents, college students, and members of Congress. “Nuclear is obsolete,” a Steve Jobs-like character tells an enthusiastic audience at the Slaughterbot’s product launch. “Take out your entire enemy, virtually risk-free.”

At the end of the video, which has been viewed more than three million times on YouTube, the Berkeley computer scientist Stuart Russell says into the camera, “Allowing machines to choose to kill humans will be devastating to our security and freedom.” Russell is among a group of prominent academics and tech executives, including Musk, Stephen Hawking, and Noam Chomsky, who signed on to a letter calling for a ban on “offensive autonomous weapons beyond meaningful human control.”

And yet artificial intelligence is already driving a worldwide arms race. In 2020, global spending for military A.I. was estimated to exceed six billion dollars, and is expected to nearly double by 2025. Russia is developing unmanned vehicles, including robotic tanks and surveillance systems. Last year, it was reported that Libya launched an autonomous drone that appeared to be equipped with “real-time image processing,” to identify and kill enemy fighters. Robert Work, the former Deputy Secretary of Defense, told me that intelligence suggests that China has turned decommissioned fighter jets into autonomous suicide drones that can operate together as a swarm. "That becomes an entirely new kind of weapon that's extraordinarily difficult to defend against," he said.

The United States, too, is testing the use of swarming drones. In an experiment last April, a drone swarm attacked a naval vessel off the coast of California. In October, the Skyborg program, an Air Force project to build autonomous aircraft to serve alongside F-35 pilots, tested two drones in live flight. Skyborg drones will be able to detect ground and air threats, identify suitable "kill" targets, and aim weapons for an optimal strike. The actual decision to "employ lethality," as the Air Force calls it, will remain in the hands of a human pilot. But, in 2020, the Air Force's chief scientist, Richard Joseph, cautioned that "we have some other questions to answer. How much autonomy do we want for a system that can deliver lethal force, and especially one that's moving at machine speed?"

"Have you tried turning it off and taking a nap?"

Cartoon by Carolita Johnson

In a paper published last April, Robert Work wrote that A.I.-enabled systems "are likely to help mitigate the biggest cause of unintended combat engagements: target misidentification." The U.S. military has repeatedly promised that improved technology would enhance enemy targeting. The results have been mixed. In 2003, during the Iraq War, an early autonomous weapon, the Patriot missile, inadvertently shot down a British fighter jet, killing both pilots, and a Navy plane, killing that pilot as well. A subsequent Pentagon report concluded that the human operators had given too much autonomy to the missile system. In a recent examination of thirteen hundred classified reports of civilian casualties in the Middle East, the *Times* characterized the American air war as "a sharp contrast to the American government's image of war waged by all-seeing drones and precision bombs."

Pettyjohn, of the Center for a New American Security, told me that the military is currently developing autonomous systems to help identify targets. "And that's one of the things that A.I. still struggles with," she said. "It's still a really hard thing to do—discriminate in the air when you're ten or twenty or thirty thousand feet in the sky." In 2018, researchers at M.I.T. and Stanford found that three popular A.I. facial-recognition systems often failed to identify the gender of women with dark skin. Two years later, a Congressional Research Service report noted that "this could hold significant implications for A.I. applications in a military context, particularly if such biases remain undetected and are incorporated into systems with lethal effects."

Stop Killer Robots, a coalition of more than a hundred and eighty non-governmental organizations, including Amnesty International, Human Rights Watch, and the World Council of Churches, has urged nations to adopt a legal treaty controlling the use of lethal autonomous weapons. The U.S. is not among the nearly seventy countries that have so far signed on. "It's not just about banning a particular weapon, like we ban land mines or chemical weapons," Bonnie Docherty, a lecturer on human-rights law at Harvard Law School and a senior researcher in the arms division at Human Rights Watch, said. "This is an effort to preëmpt the development of a technology that could alter the way wars are fought in a really dreadful way."

The D.O.D.'s position on lethal autonomous weapons, established in 2012, requires a human decision-maker to remain in the loop to an "appropriate" degree. David Ochmanek, a senior defense analyst at the Rand Corporation, whose former office at the Pentagon drafted the 2012 directive, told me, "It does not, in fact, prohibit the development of autonomous weapons." Rather, he added, "it puts in place a number of processes for review and safeguards. The commander has to be able to intervene and turn on the autonomy and turn it off as needed."

Ochmanek sees the development of autonomous weapons as a matter of deterrence, particularly against large-scale acts of aggression, such as Russia invading *NATO* territory or China invading Taiwan. “Can autonomy in different manifestations enable us to credibly believe we could defeat an invasion of this type?” he said. “The answer to that question is very much yes.”

In Niagara Falls last spring, as the L-39 was flying over Lake Ontario, the *ACE* program’s scientists and engineers gathered for their regular quarterly meeting. In a session with the groups competing to design the dogfighting algorithm, Chris Gentile, of EpiSci, emphasized that the program was not creating lethal autonomous weapons: “What we are doing is building tools to enable pilots to execute decisions more effectively.” But, as artificial intelligence ramps up the speed of decision-making, the question ultimately may be why have a human in the cockpit at all. “The Defense Department will tell you that they’re not going to have totally autonomous systems,” Pettyjohn told me. “But I have a hard time imagining, when everything is premised on making decisions faster than your adversary does, how people can actually be in that loop.”

In late September, I observed another trust experiment at O.P.L. For much of the day, a veteran pilot sat in the cockpit of one of the L-29s, which Schnell had turned into a flight simulator. Wearing a V.R. headset, he performed the battle-management role as the A.I. fought a series of scrimmages. Like the young pilot I’d watched the previous spring, he was asked to rate his trust in the A.I. while his biometric and flight data were recorded. But this time, instead of using pre-scripted scenarios, he was dogfighting with the three A.I. agents that had survived elimination: those developed by Heron Systems, PhysicsAI, and EpiSci.

Toward the end of the day, the parameters of the experiment were changed. The pilot was allowed to paddle off, fly the plane manually, then, when he felt that it was safe, cede the controls back to the A.I. According to Lauren Reinerman-Jones, a senior scientist at SoarTech, the researchers’ expectation was that, if the pilot lost the first scrimmage, it would take several more to recover trust. But, if he won the first fight, his trust in the A.I. would carry over to subsequent scenarios. Then, if he lost the final scrimmage, trust would decrease, but to a lesser degree—in coffee cups rather than buckets.

Four computers were positioned next to the plane. One recorded what the pilot was seeing in his headset. Another graphed his physiological responses, which were translated into various types of trust processing, each represented by a different-colored line. Reinerman-Jones explained that a brown line, displayed above the rest, aggregated the data into a crude rendering of the trust model she and her colleagues were developing. Woodruff sat nearby, with a computer on her lap and a recorder in hand. Every minute or so, she'd ask the pilot to assess his trust in the A.I. Almost invariably, he said it was high.

But during his debriefing he expressed some frustration with the experiment. In one scrimmage, his plane and the adversary's chased each other around and around—on the screen, it looked like water circling a drain. The pilot told Woodruff that, though he let the A.I. keep fighting, "I know it is not going to gun this guy anytime soon. In real life, if you keep going around like that you're either going to run out of gas or another bad guy will come up from behind and kill you." In an actual battle, he would have accepted more risk in order to get to a better offensive angle. "I mean, A.I. should be so much smarter than me," he said. "So if I'm looking out there, thinking I could have gained some advantage here and A.I. isn't, I have to start asking why."

For the moment, the pilot's critique reflects the immaturity of the A.I. agents, which will need much more training if they are to become sophisticated enough to take on a real adversary. But it also harks back to what Justin Mock said at the AlphaDogfight Trials a year earlier: fighter pilots trust what works.

Since then, the teams had been developing and testing A.I. agents that could take on two adversaries simultaneously, a far more complicated task. They were also beginning to develop the tools that would enable them to advance to live flight. The first step was to integrate the L-39's flight dynamics into the dogfighting algorithms and test them in drones. In the next few months, the program will be putting the A.I. agents into an airborne simulator, so that the pilots can experience g-forces, and see how that affects trust in A.I. "There's a saying in the flight-test community: All models are wrong, some models are useful," Ryan Hefron, the *ACE* program's manager, said. "So we have to find the useful pieces."

When I discussed some of these advances with Schnell, he said, "Everyone's a hero in the sim." The stakes are easy to overlook inside the O.P.L.'s hangar. Nobody gets hurt or killed crashing a virtual plane. "To truly trigger this trust equation we're working on," he told me, "you have to have another piece of metal coming right at you." ♦

An earlier version of this article misquoted Paul Schifferle.