

Q&A: Missy Cummings strives to make self-driving cars safer

The former fighter pilot wants the designers of autonomous systems to heed the hard lessons from aviation.

By **Jenessa Duncombe**



▲ Missy Cummings (Photo courtesy of George Mason University.)

Mary “Missy” Cummings became a US Navy fighter pilot in the 1980s at a turning point in aviation. New technology was transforming analog aircraft into computer-driven machines.

Over her nine years as a pilot, Cummings mastered how to fly without, and then with, a computer’s help. She witnessed military aircraft accidents caused by human-machine miscommunication. “You’re not sure who’s in control,” she says. “Is it you? Is it the computer? What are you supposed to do?” During her time as a pilot, high-profile crashes from human-autopilot confusion in commercial aviation led to federal investigations. The consequences of the disconnect between human and machine inspired Cummings to go back to school for her PhD in systems engineering.

Now the director of George Mason University’s autonomy and robotics center, Cummings studies how the designers of self-driving cars and other systems can avoid accidents when introducing automation. “A lot of people had to die in the ’80s and ’90s for us to learn these lessons,” says Cummings. Her research applies a human-centered design framework to new autonomous systems.

The following interview has been edited for length and clarity.

What motivated you to become a fighter pilot?

The Berlin Wall had not yet come down, and *Top Gun* had just come out. In my mind, why would you not want to be a fighter pilot? When I found out when I was in the Navy that women could be pilots, then there was no other path.

How did planes change during your service?

In fly-by-wire aircraft, like F-18s, the planes are doing all the flying, meaning the computer commands the plane's ailerons and rudder to go to certain positions. Where you put the plane's center stick is translated by the computer into actually controlling the air surfaces. Whereas in the analog A-4 aircraft that I flew first, I had direct control over the ailerons and rudder.

What were the implications of adding automation?

In the three years that I flew F-18s, one person a month died on average, and it was all human-machine interaction problems. In aviation, that was the first time we took a safety-critical system and put a lot of software in it. Some commercial planes flew into the side of mountains or hit other aircraft.

There's a famous graph of air-

craft accidents. Aviation had just been getting safer over time in general. This little bump of accidents in the '80s and '90s happened while I was in the military flying, and that's what motivated me to go back to school.

Your research looks at the cognitive load on humans when they interact with machines. It's tempting to think that adding more digital information to environments like a car dashboard would help people make better decisions. But your research often finds the opposite is true.

Not enough information can be detrimental, but too many streams of information, especially if you don't know how to prioritize that incoming information, can be harmful. And that's where we find ourselves. If you look at modern-day cockpits, the number of dials, sensors, gauges—you can very quickly overwhelm a pilot by pushing in a lot of information and alarms. For example, there was a case where an Airbus in 2010 en route from Singapore to Sydney had a major engine failure, and it sounded like a Las Vegas showroom in the cockpit because the alarms hadn't been coordinated, and the pilots were just being

bombarded on all sides.


The problem with people having more information in cars is that it diverts their attention from the highly visual task of driving, so there is a balance needed. We want to give people enough information to do their jobs but not too much and overwhelm them, like in the Airbus incident. Or deliver irrelevant, distracting information, like when BMW pushed *Spider-Man* ads on drivers' dashboard screens this year.

We humans are in love with an inflow of information, but unless we really take a human-centered design approach, we can actually make people's jobs much harder, if not impossible.

As a researcher, how do you investigate that?

I've spent basically a decade of my research career showing people that you need to be able to construct a computational model of what information is coming in, what the capabilities are of the people, and what the task demands are from the environment. Then you can dynamically model pretty well when and how you can task-load people.

For instance, we looked at how people's attention waned when driving for long periods. We put





▲ As a US Navy fighter pilot, Cummings felt the computer in early automated aircraft hiccup at times. (Photo courtesy of Missy Cummings.)



▲ Can a car's autonomous alert system spot an inflatable test dummy crossing the street? Cummings put several recent car models to the test as part of her academic research on autonomous systems. (Photo courtesy of Missy Cummings.)

people behind the wheel in a simulated driving scenario and measured their brain activity for four hours. Even when they were in charge of driving, people's brain activity dropped on average after about 20 minutes. Self-driving cars that rely on driver assistance to intervene in unsafe situations will need to find ways to detect how much people are paying attention to know whether a human could really intervene in time.

What other issues do self-driving cars face?

There is no such thing as truly self-driving cars, meaning that all the self-driving cars today have remote operators that are babysitting them. My research on dispatch operations for self-driving cars found that the division of tasks is important. You don't want to over- or underburden someone. In December, there was a power outage in San Francisco and Waymo's self-driving cars came to a halt, putting the city in gridlock, because they were waiting to hear from their remote operators.

What about the safety of automated driver assists?

There are two levels of driver assist. One is a must-have, one is a nice-to-have. The must-haves are safety driver-assist features, like blind-spot warnings or automatic braking. Those are must-haves because they are all passive. They sit in the background, and they help you be safer.

The nice-to-haves are different. Automated driver assist passes cars for you or does automated cruise control. It controls your speed and takes exits.

My issues with driver assist are with the convenience features, not the safety features. We recently tested whether cars in automatic driving mode, in which no driver has hands on the steering wheel, could reliably detect a pedestrian. The cars drove on a test track toward an inflatable dummy at 40 miles per hour. We repeated the test a couple hundred times. Our research shows that one driver-assist system on the market today cannot reliably detect a pedestrian.

What role could generative AI have in automation?

When you play around with a large language model, you can type in a question and then get an answer. If you type in the same question again, you'll get a different answer. Neural nets are nondeterministic by design, which is fine. And indeed, that's what's given us some creative abilities in generative AI, which can be very helpful. But this is why it's two sides of the same coin. You can get creativity out of uncertainty, but you can also get unpredictable performance. The probability that you will have an incorrect answer significantly grows in ways that are not appreciated by the medical, transportation, and defense communities.

There are discussions of putting large language models in charge of controlling weapons. There has never been a worse idea. I even hear questions about putting these technologies inside of nuclear reactors. No, no, no, no, no! Academia and industry have fallen in love with generative AI but don't want to be forced to deal with its messiness, and that messiness is uncertainty. **PT**