The Economist

Aviation safety

Flight response

An artificially intelligent autopilot that learns by example

Sep 17th 2016 | From the print edition

ON JUNE 1st 2009, an Air France airliner travelling from Rio de Janeiro to Paris flew into a mid-Atlantic storm. Ice began forming in the sensors used by the aircraft to measure its airspeed, depriving the autopilot of that vital data. So, by design, the machine switched itself off and ceded control to the pilots. Without knowing their speed, and with no horizon visible in a storm in the dead of night, the crew struggled to cope. Against all their training, they kept the plane's nose pointed upward, forcing it to lose speed and lift. Shortly afterwards the aeroplane plummeted into the ocean, killing all 228 people on board.

Just relax and enjoy the view, Captain

French air-accident investigators concluded that a lack of pilot training played a big part in the tragedy. As cockpits become ever more computerised, pilots need to keep their flying skills up to date. But pilots are also in short supply. In July Airbus predicted that 500,000 more will be needed by 2035 to keep pace with aviation's expected growth. That means there is pressure to keep aircrew in their cockpits, earning money, rather than in the simulators, taking expensive refresher courses.

Help may be at hand, though, from artificial-intelligence (AI) experts at University College London (UCL). Inspired by the Air France tragedy, Haitham Baomar and his colleague Peter Bentley are developing a special kind of autopilot: one that uses a "machine learning" system to cope when the going gets tough, rather than ceding control to the crew.

Today's autopilots cannot be trained, says Mr Baomar, because they are "hard coded" programs

in which a limited number of situations activate well-defined, pre-written coping strategies—to maintain a certain speed or altitude, say. A list of bullet points (which is what such programs amount to) does not handle novelty well: throw a situation at the computer that its programmers have not foreseen, and it has no option but to defer to the humans.

Mr Baomar suspected that a machine-learning algorithm could learn from how human pilots cope with serious emergencies like sudden turbulence, engine failures, or even—as happened to the Air France jet—the loss of critical flight data. That way, he says, the autopilot might not have to cede control as often, and that, in turn, might save lives.

AI takes off

Machine learning is a hot topic in AI research. It is already used for tasks as diverse as decoding human speech, image recognition or deciding which adverts to show web users. The programs work by using artificial neural networks (ANNs), which are loosely inspired by biological brains, to crunch huge quantities of data, looking for patterns and extracting rules that make them more efficient at whatever task they have been set. That allows the computers to teach themselves rules of thumb that human programmers would otherwise have to try to write explicitly in computer code, a notoriously difficult task.

UCL has lots of experience in this area. It was the institution that spawned DeepMind, the company (now owned by Google) whose AlphaGo system this year beat a human grandmaster at Go, a fiendishly complicated board game. The UCL team has written what it calls an Intelligent Autopilot System that uses ten separate ANNs. Each is tasked with learning the best settings for different controls (the throttle, ailerons, elevators and so on) in a variety of different conditions. Hundreds of ANNs would probably be needed to cope with a real aircraft, says Dr Bentley. But ten is enough to check whether the idea is fundamentally a sound one.

To train the autopilot, its ten ANNs observe humans using a flight simulator. As the plane is flown—taking off, cruising, landing and coping with severe weather and aircraft faults that can strike at any point—the networks teach themselves how each specific element of powered flight relates to all the others. When the system is given a simulated aircraft of its own, it will thus know how to alter the plane's controls to keep it flying as straight and level as possible, come what may.

In a demonstration at a UCL lab, the system recovered with aplomb from all sorts of in-flight mishaps, from losing engine power to extreme turbulence or blinding hail. If it were to lose speed data as the Air France flight did, says Mr Baomar, the machine would keep the nose low enough to prevent a stall. The newest version will seek speed data from other sources, like the global positioning system (GPS).

To the team's surprise, the system could also fly aircraft it had not been trained on. Despite learning on a (simulated) Cirrus light aircraft, the machine proved adept with the airliners and fighter jets also available in the database. That is a good example of a machine-learning phenomenon called "generalisation", in which neural networks can handle scenarios that are conceptually similar, but different in the specifics, to the ones they are trained on.

UCL is not the only institution interested in better autopilots. Andrew Anderson of Airbus, a big European maker of jets, says his firm is investigating neural networks, too. But such systems are unlikely to be flying passenger jets just yet. One of the downsides of having a computer train itself is that the result is a black box. Neural networks learn by modifying the strength of the connections between their simulated neurons. The exact strengths they end up with are not programmed by engineers, and it may not be clear to outside observers what function a specific neuron is serving. That means that ANNs cannot yet be validated by aviation authorities, says Peter Ladkin, a safety expert at Bielefeld University in Germany.

Instead, the new autopilot will probably find its first uses in drones. The system's versatility has already impressed delegates at the 2016 International Conference on Unmanned Aircraft Systems in Virginia, where Mr Baomar presented a paper. The system's ability to keep control in challenging weather might see it used in scientific investigations of things like hurricanes and tornadoes, says Dr Ladkin—some of the most challenging flying there is.

From the print edition: Science and technology