



Why Aviation Is the Hardest Physical AI Problem

April 2, 2026

By Matt George, Founder & CEO, Merlin Labs

There is a version of the Physical AI story that sounds straightforward: train a model in simulation, deploy it in the real world, iterate. Companies are doing this for warehouse floors, urban roads, and factory assembly lines. This is complex work, requiring a lot of engineering and compute power, as well as significant applied experience on the part of really innovative teams.

But autonomous aviation is a different problem. A warehouse robot navigating a flat floor with fixed landmarks is operating in a constrained, structured environment. The physical variables are manageable. Failure modes are recoverable. A self-driving car operates in two dimensions, at relatively low speeds, with well-mapped road geometry and decades of infrastructure designed around human reaction time. Put another way, if needed, it's relatively straightforward for a self-driving car to pull over. Landing a plane if something isn't right is not so simple.

Merlin Pilot operates in three-dimensional space, at 300+ knots, in conditions that change continuously and often without warning. The physical variables are not just information to be managed. They are, in many cases, the mission itself. I.e., weather, turbulence, aerodynamics, propulsion systems, structural load limits, airspace deconfliction, mechanical failure modes: these all represent an aircraft's status in its operating environment, doing whatever it is supposed to be doing - they are not just edge cases.

The certification bar reflects that.

To deploy autonomy on a commercial or military aircraft, you do not demo it. You certify it, to aviation standards, through a process that takes years and requires the system to perform reliably across the full envelope of the aircraft's operational life. The New Zealand Civil Aviation Authority and the U.S. Air Force do not issue airworthiness approvals because a system worked in test conditions. They issue them because the system works, consistently, in the conditions that could actually kill people when things go wrong. These are the certification authorities and regulators with whom Merlin is already working.

This is the point most Physical AI narratives skip past. Simulation fidelity, training data and architecture all matter. But in another sense, none of that matters *in aviation* until autonomy has cleared a regulatory process that treats failure as unacceptable, rather than a learning event. Certification is not just a process gate. It's proof the system works reliably enough in the physical world to be trusted with an aircraft.

Each Merlin Pilot autonomous flight generates physics data from the real environment: the actual aerodynamic response of a specific airframe in actual weather, at actual altitude, on actual missions. That data feeds back into the system. Every flight expands the training set. The more we fly, the more the system understands about how the physical world behaves in conditions that matter.

That loop, i.e., real-world operation generating data that compounds system intelligence, is what Jensen Huang describes when he talks about Physical AI at scale. Merlin is already inside that loop.

The economic implications of all of this are straightforward. Aviation arguably is the most complex Physical AI domain. The safety requirements are the most demanding. The regulatory bar is the highest. And we believe the economic value per deployed unit dwarfs what the near-term consumer robotics market generates per unit. In this context, the difficulty of the problem is not just a cost, it is a moat.

We are not claiming to be part of the Physical AI wave because the category is attracting capital. We are making the case that aviation is the hardest version of the problem and that we believe Merlin is further along in solving it than any other company in the field.

This post contains forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995. All statements other than statements of historical fact should be considered forward-looking statements, including statements regarding Merlin Labs' competitive position and relative standing in the autonomous aviation market, the anticipated economic value of autonomous aviation relative to other physical AI domains, the future performance and continued development of Merlin Pilot, and anticipated regulatory certifications and approvals. These forward-looking statements are based largely on our current expectations and projections about future events and trends. These forward-looking statements are neither promises nor guarantees, but involve known and unknown risks and uncertainties that could cause actual results to differ materially from any future results, performance or achievements expressed or implied by the forward-looking statements, including without limitation the risks, uncertainties, and assumptions described under "Risk Factors" in Merlin Labs' S-4/A filed with the SEC on February 9, 2026. Except as required by applicable law, we undertake no obligation to update any of these forward-looking statements for any reason after the date of this post.