LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

Why Can't AI Fly a Plane? Technical & Human Barriers

By landing.aero   Published January 21, 2026   41 min read

Executive Summary

The idea of fully autonomous, pilotless commercial airliners is no longer science fiction, but several formidable hurdles still stand in the way. Decades of incremental automation have given today's large transport aircraft highly capable autopilot and flight-management systems – modern jets routinely fly extended periods on autopilot under human supervision. Nevertheless, a fully "AI pilot" remains unachievable in practice for now. **Technically**, although prototypes and experiments (e.g., Airbus's image-guided **ATTOL** tests (Source: aerospaceamerica.aiaa.org) (Source: www.mdpi.com) and NASA's **ICAROUS** autonomy suites (Source: www.mdpi.com) (Source: www.mdpi.com) demonstrate that software can taxi, take off, and land without human hands on the controls, these systems have not been proven to the extremely high reliability demanded in commercial aviation. The airline industry's "zero-failure" safety culture requires on the order of 10^{-9} probability of fatal accident per flight hour (Source: www.mdpi.com) (Source: www.mdpi.com), a standard that learning-based AI has not yet demonstrably met. Machine-learning algorithms are inherently **nondeterministic** – they may behave differently even in identical situations – which conflicts with existing certification standards (e.g. **FAA DO-178C**) designed around deterministic, hand-coded software (Source: aerospaceamerica.aiaa.org) (Source: www.mdpi.com). In short, **technical verification** of AI-based flight control poses unresolved challenges.

Equally important are **human and institutional factors**. Surveys of passengers and pilots alike show deep unease with pilotless commercial flights: 72% of passengers cite *safety* and 60% *human judgment* as concerns, and 80% of pilots oppose removing the human pilot entirely (Source: www.mdpi.com) (Source: www.mdpi.com). This lack of trust reflects both fear of the unknown and recognition that AI lacks the common-sense and intuition that human pilots apply in crises (Source: www.mdpi.com) (Source: www.mdpi.com). Trust-building efforts (e.g. "human-on-the-loop" cockpit concepts and explainable-AI) may help, but public acceptance remains far from guaranteed.

Finally, **regulatory and economic barriers** loom large. No current aviation regulation permits fully autonomous passenger flights; authorities such as FAA, EASA, and ICAO are only beginning to draft frameworks for learning-enabled aircraft (Source: www.mdpi.com) (Source: aerospaceamerica.aiaa.org). Certification of an AI pilot would require entirely new processes, test regimes, and possibly frozen ("weight-frozen") models or continuous re-validation (Source: www.mdpi.com) (Source: aerospaceamerica.aiaa.org). Even if technically viable, the costs of certification, ground infrastructure (secure data links, control centers), and liability protection may outweigh the modest operational savings (on the order of 10–

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

15% from eliminating crew) (Source: www.mdpi.com). In sum, while the *principles* of autonomous flight are coming into focus, the **reality** is that a confluence of technical, human, and regulatory factors must be resolved before AI can "take the pilot seat." In the near term, experts recommend incremental steps (autonomous cargo flights, one-pilot operations, advanced autopilots) rather than immediate pilotless passenger service (Source: www.mdpi.com) (Source: aerospaceamerica.aiaa.org).

Introduction and Background

The last century of aviation has seen a steady march of automation, but **fully AI-driven flight remains out of reach**. Early autopilots appeared in the 1930s, allowing aircraft to fly straight-and-level without constant pilot inputs. Modern airliners built in the 21st century routinely use sophisticated Autopilot/Flight Management Systems (FMS) that can fly nearly every phase of flight when guided by a human pilot. Today's jets can take off, climb, cruise on programmable routes, descend, and even land automatically (especially in low-visibility conditions using autoland) under human supervision. In this context, the notion of a pilot serving merely as a monitor – or eventually, not at all – is technically plausible.

However, **technical feasibility is only one side of the problem**. As Elbasyouny & Dababneh (2026) note, transitioning to *Pilotless Passenger Aircraft* (PPAs) requires "rigorous assurance of safety, certification of complex AI/automation systems, clarity on liability, and broad public trust" (Source: www.mdpi.com). That is, even if the engineering can be solved, aviation's overarching imperative is never to compromise safety. Commercial air travel today is the world's safest mode of transportation, thanks to institutionalized safety management systems, conservative design criteria, and extensive certification protocols (Source: www.mdpi.com). Introducing AI into this ecosystem upends many assumptions. Aviation's history privileges a "zero-failure" ethos (Source: www.mdpi.com): stakeholders insist that any new automation matches or exceeds existing reliability standards. The challenge of an AI pilot is thus *sociotechnical*: it transcends mere coding problems and touches on human factors, regulation, cybersecurity, and economics (Source: www.mdpi.com) (Source: www.mdpi.com).

Prior research and commentary corroborate these themes. Aerospace industry experts observe that **today's automated systems fly much of the plane**, but cities of regulators and pilots alike emphasize their limits (Source: aerospaceamerica.aiaa.org) (Source: www.wired.com). For example, Jon Kelvey's Aerospace America feature summarizes Asiana Flight 214 (2013) to illustrate how existing automation *failed* to prevent a fatal accident when pilot input was mishandled (Source: aerospaceamerica.aiaa.org). Likewise, Wired magazine (2008) observed that while autopilots "do most of the work once a plane is aloft" (Source: www.wired.com), fully removing humans from the cockpit introduces unprecedented certification and design hurdles.A recent MDPI study combining passenger surveys and pilot interviews underscores the multifaceted barriers: technical reliability, regulatory approval, human supervision, and public acceptance were repeatedly cited as prerequisites for any pilotless flight operation (Source: www.mdpi.com) (Source: www.mdpi.com).

In sum, our goal here is to examine **what is stopping AI from flying a plane**, from every angle. We consider the **technological state of the art** (AI and automation capabilities), the **regulatory landscape** (certification obstacles), the **human factor** (trust and usability), and **case studies** (notable incidents and experiments). Wherever possible, we ground claims in data, survey results, and expert analyses. By profiling the challenges in depth, we show that **autonomous aviation remains an open frontier** – even as breakthroughs in sensors, machine learning, and computation steadily advance, aviation's demands keep the finish line in sight, not within immediate reach.

Current State of Aviation Automation

Traditional Autopilot and Flight Management

Understanding why AI hasn't replaced pilots starts with grasping what automation *can* do today. Modern airliners are endowed with advanced flight management systems (FMS) and autopilots that relay pilot commands, manage navigation, and even fly ILS landings. In basic terms, "autopilot is, at the most basic level, pretty simple," notes Wired: it **uses the pilot's inputs to adjust and maintain the airplane's heading, altitude, and speed** (Source: www.wired.com). The pilot selects waypoints or altitude requirements on the FMS, and the autopilot follows those instructions by moving the control surfaces. In less capable aircraft (general aviation), a pilot might "fly by the needles" on an ILS approach, whereas airliners can plug those flight director cues into autoland software. In rough weather or low visibility, autoland often keeps jets steadier on glidepaths than humans (Source: aerospaceamerica.aiaa.org). In this sense, existing automation is very effective at routine, well-defined tasks.

However, even today's autopilots **rely fundamentally on human oversight**. A copilot or pilot is always setting the objectives and monitoring results. In fact, airlines train pilots to be vigilant: autopilots are not a free lunch. John W. Hannigan of the NTSB famously noted in the Asiana 214 investigation that "in this instance, the flight crew over-relied on automated systems they did not fully understand" (Source: aerospaceamerica.aiaa.org). Modern cockpits retain manual controls, and procedures ensure pilots are ready to intervene. For instance, if the autopilot disconnects or an abnormal

situation arises, the pilot must immediately take control. Moreover, autopilots generally handle only **nominal conditions** – stable cruise, approach under functioning ILS, etc. They are *not* programmed to handle unusual events like a multi-system failure or unforeseen obstacle; those still require human judgment (Source: aerospaceamerica.aiaa.org).

Another distinction is that autopilots are **deterministic, rule-based systems**. Every autopilot function is the result of pre-coded feedback loops (altitude-hold, yaw-damper, VNAV, etc.) which can be tested and certified to meet DO-178C standards for software. There is no "learning" or adaptiveness; given the same inputs, a traditional autopilot will always respond the same way. By contrast, **AI-driven autopilots** (based on machine learning) would determine actions through data-trained models. This nondeterministic nature underlies a major regulatory obstacle (discussed in Section 4): as Singh puts it, if an AI runs a scenario a hundred times and "40 times would go left, and 60 times go to the right," that *"nondeterminacy does not meet the [existing DO-178C] standard"* (Source: aerospaceamerica.aiaa.org). In short, while autopilots today are well-understood and certifiable, AI systems introduce unpredictability that current regs cannot easily accommodate.

Recent Demonstrations of AI Autonomy

Despite these limits, researchers and manufacturers have pushed automation well beyond the basics. Airbus, Boeing, NASA, and others have run experiments in recent years to test "beyond-autopilot" technologies. Notable examples include:

- **Airbus ATTOL (Autonomous Taxi, Take-Off & Landing)**: In 2019–2020 Airbus performed trials on a modified A350-1000. Cameras installed in the nose imaged the runway environment, and onboard image-recognition algorithms "identified features to 'see' the runway" without ILS/{{Radio Navigation}} (Source: aerospaceamerica.aiaa.org). Over two years and hundreds of test flights, the system **learned** to handle taxiing, takeoff, and landing via machine vision. These flights succeeded under supervision, demonstrating that fully automated surface operations are technically viable (Source: aerospaceamerica.aiaa.org) (Source: aerospaceamerica.aiaa.org). However, Airbus notes the system was always flown with a safety crew ready to take over, and importantly, "governments have no process in place for permitting automation such as ATTOL and IAS aboard airliners" (Source: aerospaceamerica.aiaa.org).

- **NASA ICAROUS**: NASA's Integrated Configurable Algorithms for Reliable Operation of UAS (ICAROUS) project has produced a suite of software for autonomous navigation, collision avoidance, and contingency planning in complex airspace (Source: www.mdpi.com) (Source: www.mdpi.com). While aimed originally at unmanned systems, ICAROUS embodies *safety-centric* path planning and detection algorithms. Tests show ICAROUS can let drones or aircraft autonomously navigate controlled-class airspace, provided the mission stays within a predefined envelope (Source: www.mdpi.com) (Source: www.mdpi.com). Similarly, NASA and DARPA have flown transport-category planes remotely or autonomously in demonstration flights (e.g. the DARPA "ALIAS" experiment with a Learjet), showing that large jets can loop into automated control under supervision.

- **Prototype AI Autro-Pilots**: Independent research groups have trained AI agents on flight simulators. For instance, Baomar and Bentley (2016–2021) developed an Intelligent Autopilot System (IAS) neural network by feeding it thousands of hours of Boeing 787 simulator flights. After training, IAS managed never-before-seen scenarios: for example, it held glideslope in 50–70 knot crosswinds where the standard autopilot would have disengaged (Source: aerospaceamerica.aiaa.org). Studies suggest such AI controllers could be retrofitted into cockpits, replacing conventional autopilot algorithms. However, these remain experimental and have not (yet) flown real aircraft. They do highlight that **AI can exceed current autopilot envelope** in controlled tests, but scaling that to every possible scenario in real flight is daunting.

- **Electric VTOL and Air Taxis**: The rise of electric Vertical Takeoff and Landing (eVTOL) designs for urban air mobility has spurred autonomy research. A handful of eVTOL projects (like Boeing's subsidiary Wisk Aero, Aurora Flight Sciences, etc.) are working hand-in-hand with regulators on autonomous certification. For example, Wisk aims for an autonomously piloted air taxi certification, leveraging decades of procedural knowledge to certify new systems (Source: spectrum.ieee.org) (Source: spectrum.ieee.org). These programs often start with two humans, then reduce to one and eventually aim for no onboard pilot. Currently, eVTOLs operate under very restricted conditions, often with remote pilots on hand. Still, the eVTOL experience shows industry commitment to eventual autonomy: many designs incorporate advanced sensors, flight control redundancies, and "auto-landing" modes from day one, anticipating a path to autonomy.

Table 1: Surveyed Concerns about Pilotless Flight (from Elbasyouny & Dababneh (Source: www.mdpi.com)

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

CONCERN / METRIC	PASSENGERS (N=312)	PILOTS (N=15)
Safety concerns (e.g. reliability)	72% (Source: www.mdpi.com)	–
Cybersecurity concerns	64% (Source: www.mdpi.com)	–
Need for human judgment	60% (Source: www.mdpi.com)	–
Automation improves safety	–	93% (Source: www.mdpi.com)
Oppose removal of pilots	–	80% (Source: www.mdpi.com)

Data: Percentage of survey respondents expressing each concern or position.

The results above illustrate a key point: even as technology progresses, **social acceptance lags**. Passengers overwhelmingly worry about fundamental safety, cyberattacks, and lack of human judgment. Pilots agree that automation *can* improve safety, but most refuse to give up on having a pilot on board. This cautious stance reflects aviation's conservative culture. In practice, **autonomy in flight tends to be introduced gradually** – first as decision aids, then as "copilot" systems. It is widely expected that AI will augment pilots, not replace them overnight (Source: www.mdpi.com) (Source: aerospaceamerica.aiaa.org). Both industry roadmaps and expert opinion emphasize extended human oversight (e.g. unmodified pilot in command) for years to come.

Technical Challenges for Autonomous Flight

Unifying Complexity and Uncertainty

Commercial flight spans many complex tasks and environments: navigating through congested airspace, taxiing around airport ground traffic, handling severe weather, system failures, runway detection, and more. Human pilots rely on wide situational awareness, intuition, and split-second judgment to manage the unexpected. AI systems, on the other hand, must rely entirely on their sensors and algorithms. This introduces several challenges:

- **Environmental Variability**: The real world is full of unanticipated situations. For instance, consider runway detection. At major airports the infrastructure is well known, but still can be degraded. Airbus's ATTOL trials showed that cameras and ML can recognize runways in clear conditions (Source: aerospaceamerica.aiaa.org), but poor weather, low light or obscured markings could foil vision-based systems. Similarly, GPS signals or radio-navigation (ILS) may be unavailable (as happened to Asiana 214 (Source: aerospaceamerica.aiaa.org), requiring alternate strategies. While humans can adapt (e.g. visually acquire approach lights, guesses), AI vision systems must be trained on extensive, representative data. Missing corner cases (e.g. a runway under snow) could cause a vision-based autopilot to fail to detect the runway line. In essence, the **operational design domain (ODD)** of an AI autopilot – the exact conditions and scenarios it can handle – must be explicitly defined and strictly maintained. Outside that envelope, catastrophic failures may occur.

- **Sensor Reliability and Redundancy**: All avionics rely on sensors (altimeters, air data computers, inertial units, radars, cameras, etc.). These sensors can fail or give misleading readings. Traditional pilots double- or triple-check every parameter. An AI system must either incorporate multiple redundant sensors (and fuse them) or have abort criteria. But designing such robust sensor fusion is nontrivial. For example, if heavy rain obscures a camera view and simultaneously an altimeter glitch occurs, can the AI reliably continue approach? Human pilots might spot other cues (PAPI lights, runway edge) or execute a go-around. An AI might misinterpret the data and continue unwisely. Missing or faulty sensor data can trigger safe-mode or emergency protocols – but these must themselves be validated.

- **Dynamic Failures and Emergencies**: One of aviation's regulatory principles is that systems must handle all *foreseeable* failures. In manned flight, the pilot is the last line of defense. In an autonomous mode, the AI would have to take that role. This includes engine failures, control surface jams, or even catastrophic structural damage. These are *rare* but **must** be managed flawlessly. Today's autopilots do not claim to address these ultra-rare events; rather, pilots explicitly train for them. For an AI to assume the pilot's role, it must have strategies for *every* emergency: multiple engine-out glide envelopes, asymmetric thrust compensation, partial control loss, etc. Implementing and verifying all of that in software is a Herculean task. In accidents like Qantas 32 (see Case Study below), automation partially failed mid-flight and humans had to hand-fly the rest. (Source: aerospaceamerica.aiaa.org) An AI pilot would need to do the same – but to allow that, the AI itself must demonstrate it can reliably recover from such emergencies *without human intervention*, an expectation far beyond current capabilities.

- **Algorithmic Predictability**: Modern machine learning (ML) systems, especially deep neural networks, can pose unpredictable behaviors. Unlike traditional code where every branch can be traced, a neural network's decision boundary is opaque. This lack of "explainability" is problematic in safety-critical domains. Regulators and manufacturers worry that an AI pilot might make unanticipated choices in corner cases. Researchers are developing techniques (e.g. formal verification of neural nets, safety wrappers) but these are in infancy. For example, if an AI pilot tries to avoid another plane by a novel maneuver not seen in training, how can we certify that attempt? It could inadvertently lead to a stall or enter restricted airspace without clear reason. Such uncertainty must be addressed by either limiting autonomy (only allow in well-tested conditions) or by adding layers of certifiable logic (e.g. safety constraints that override the AI if it attempts something outside known safe bounds).

Software and Certification Issues

The certifiability of learning-based flight software is perhaps the central sticking point. Current certification (FAA/EASA DO-178C level A/B for flight-critical systems) assumes determinism and exhaustive testing of all code paths. Neural networks violate that assumption. The Aerospace America article clarifies: "FAA's guideline (DO-178C) isn't designed to deal with neural networks that are nondeterministic, meaning they react differently to the same situation at different times" (Source: aerospaceamerica.aiaa.org). The example given is stark: if an autonomous algorithm's response to, say, an obstacle-on-runway scenario is not 100% repeatable (left half the time, right half the time), then by definition it fails DO-178C's criteria (Source: aerospaceamerica.aiaa.org). Achieving 100% determinism with a learning agent may require "gold-plating" (excessive over-engineering) or freezing a trained model in place, which itself is an open research issue.

Moreover, many certification requirements rely on having a pilot to handle failure cases. If that role is removed, the system itself must satisfy safety that typically counted on the pilot margin. Wisk Aero's work illustrates this: even for an air taxi, its FAA certification team notes "we still do our safety assessment the same way … but it may be to higher levels [of reliability]" (Source: spectrum.ieee.org). In other words, every failure mode must be made even more unlikely if no pilot is there to intervene. This drives the need for ultra-redundant architectures (multiple independent sensors, independent flight computer channels, etc.) (Source: www.mdpi.com). But the more complexity in redundancy, the harder it is to integrate, maintain, and certify.

Continuous learning is another complication. Conventional certifiable avionics do *not* change behavior post-certification. If an AI autopilot were to continue learning in the field, its behavior could drift unpredictably. Hence, regulators are likely to demand that any certified AI use a *frozen* model (no learning online) or require a re-certification after software updates – akin to how software bugs require patching and retesting today (Source: www.mdpi.com). Sitting on a static neural net may feel like an odd compromise: it gives up adaptiveness in exchange for certifiability. But it may be the only feasible path in the near term. Thus we see a bifurcation: *research prototypes* push continually updated AI, while *certifiable designs* may be high-assurance, static inference engines, at least initially.

Table 2 summarizes key technical differences between today's deterministic autopilots and prospective AI-based flight control systems:

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

ASPECT	TRADITIONAL AUTOPILOT	AI-BASED FLIGHT CONTROL
Design Philosophy	Fixed, hand-crafted control laws determined by engineers (Source: www.wired.com). Pilot inputs set targets (altitude, heading, etc.).	Learned behaviors from data (neural networks or other AI models); aims to replicate or improve pilot-like judgments.
Predictability	Fully deterministic given same inputs and state. Every mode and response can be traced through code.	Stochastic/Probabilistic. Outputs can vary due to learning weights; "black box" nature. Challenging to predict in unseen conditions (Source: aerospaceamerica.aiaa.org).
Certification Path	Mature. DO-178C process demands exhaustive testing and analysis. Deterministic nature fits existing standards (Source: aerospaceamerica.aiaa.org). Safety cases assume pilot backup.	Uncertain. No existing "part XX" for learning systems. Continuous learning or nondeterminism conflicts with DO-178C as-is. Possible solutions: freeze-model certification, periodic re-evaluation (Source: www.mdpi.com).
Computational Requirements	Typically low. Control laws are simple arithmetic and logic, easily run in flight-qualified hardware.	Potentially high. Neural networks, computer vision, or complex planning algorithms require powerful CPUs/GPUs or specialized AI chips, raising issues of weight, power, and certification of the hardware too.
Scope of Capability	Limited to well-defined tasks (follow prescribed lateral/vertical profiles, autoland, etc.). Handovers to pilot for anything outside design envelope.	Broad/agile. Can theoretically handle complex tasks (vision-based taxi, adaptive flight in turbulence) if trained. But each new capability must be taught. Risk of unexpected emergent behaviors.
Transparency	High. Pilot and engineers can usually trace how decisions are made via rule-based logic. Data logs (flight data recorders) clearly show steps.	Low. Decisions internal to neural nets are not easily interpretable. Explainable-AI research is needed to audit decisions for trust (Source: www.mdpi.com).
Redundancy Design	Well-understood methods of sensor/actuator redundancy. Secondary autopilots can takeover.	Uncharted. Would need redundant neural nets or parallel AI pipelines. Failover must be designed from first principles.
Failure Modes	Failures typically lead to lose-of-command scenarios requiring pilot action (e.g. autopilot disconnect).	Failures could be silent and subtle. Hard to specify how "unsafe learning" states would manifest.

This comparison highlights that AI-based systems, while having potential to **expand automation's envelope**, also introduce new categories of uncertainty. Deterministic autopilots operate within thoroughly bounded regimes; AI autopilots must be constrained and verified in ways not needed before.

Data and Training Limitations

Truly robust AI flight control would require training on *massive* datasets covering every conceivable flight condition and failure. In practice, such data are sparse. We have abundant data for routine flight (millions of flight hours), but far fewer examples of rare anomalies. Simulators can augment with synthetic edge-cases (stalls, multiple failures, extreme weather), but model fidelity and coverage remain imperfect. One risk is **overfitting to normal operations**: an AI trained mostly on stable flights might behave unpredictably in an emergency. Engineers are exploring "adversarial training" where the AI is exposed to intentionally challenging scenarios, but this is an unsolved field. Ultimately, verifying that an AI has been sufficiently trained (and won't fail unexpectedly) is itself a hurdle – a form of "AI assurance" problem recent research highlights (Source: www.mdpi.com) (Source: www.mdpi.com).

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

Cybersecurity and Integrity

Any AI pilot will rely on data links, sensors, and perhaps even external networks (for updates or guidance). This raises cybersecurity concerns. While airline systems have long been guarded, increasing connectivity (e.g. in-flight Wi-Fi, datalink communications) means new attack vectors. A malicious actor might try to spoof sensors (bad GPS signals, falsified weather data) or infiltrate the AI pre-flight (tainting a model update). Aerospace defense studies caution that networked AI could be vulnerable to "neural Trojan horses" or adversarial examples (Source: www.numberanalytics.com). In safety analysis terms, cyber threats become part of the system hazards, requiring defense-in-depth. Industry awareness is growing: the FAA is already promulgating new cybersecurity requirements for wired architectures (Source: www.faa.gov).

While no credible evidence yet suggests hijacking a large airliner via cyberattack (the FAA says stuffing 747 flight control is infeasible (Source: www.msspalert.com), **autonomy multiplies attack surfaces**. For instance, a pilotless cargo plane might be less tempting to hijack in-cockpit, but if it relies on an open internet link, intercepting or taking control remotely is conceivable. Thus, deploying AI pilots will likely demand unprecedented cybersecurity assurance: layered firewalls, anomaly detection, and possibly even cryptographic redteaming. These add cost and complexity. Indeed, some analysis suggests the new risk-mitigation measures (secure datalinks, hardened processors, redundancy against jamming) could wipe out the operational cost savings of autonomy (Source: www.mdpi.com).

Human Factors and Trust in Autonomous Flight

Even technically flawless automation can fail socially if humans cannot trust it. Trust in automation is a well-studied area: humans tend to either **over-trust** (to the point of complacency) or **under-trust** (refusal to use helpful features), depending on how systems are presented (Source: www.mdpi.com). Surveys in the aviation domain show many passengers fall into the latter camp for uncrewed flights. A small-scale experiment recorded how fliers reacted when told "this flight to New York is operated entirely by a computer with no human pilot." The result was shock: "confusion, concern and outright fear" immediately filled the gate (Source: www.aerosociety.com). Participants conflated aviation with still-developing autonomous cars ("They still have accidents"), highlighting a pattern: *if autonomous cars frighten people, autonomous planes do even more so* (Source: www.aerosociety.com).

Aircraft designers are acutely aware of this trust gap. Automated-vehicle studies show that **demonstrable safety case** and gradual exposure are key to building trust (Source: www.mdpi.com). In aviation, the concept of "adaptive automation" has arisen: AI systems that **interact** with pilots, allowing them to stay engaged. For example, an AI copilot might suggest alternate routes, or perform maneuvers while clearly communicating actions ("Pilots see this line-of-sight path over the terrain, and updated guidance on their displays"). Explainable AI is a related idea: if the system can articulate *why* it chose a maneuver ("I turned left to avoid storm cells ahead"), pilots and crew may feel more comfortable. The MDPI study explicitly recommends "explainable-AI integration" to bridge trust gaps .

Pilots themselves also have cognitive concerns. When automation handles more flying, pilots' manual skills can atrophy. Numerous studies (cited by Elbasyouny & Dababneh (Source: www.mdpi.com) document how heavy reliance on autopilot leads to degraded situational awareness: pilots may become "out of the loop," unsure what the automation is doing or why. This was a factor in Asiana 214: once the autopilot disengaged unexpectedly, the pilot was caught off-guard and made incorrect manual inputs (Source: aerospaceamerica.aiaa.org). An AI autopilot that works "too well" could ironically exacerbate this: the better it is, the more pilots fly hands-off, and the less practice they get for emergencies. Capt. Clint Balog of Embry-Riddle warns of an "automation surprise" effect: "the more [the autopilot] does, the less transparent it becomes. And the more difficult it becomes to keep the pilot in the loop when the automation fails and the pilot has to take over." (Source: aerospaceamerica.aiaa.org)

Trust variation is also cultural. Passengers often anthropomorphize risk: research suggests people **forgive human error more readily** than machine error (Source: www.mdpi.com). In a collision or crash, the notion that "a computer made this choice" can trigger a deep lack of acceptance, even if the outcome is identical. Ethics scholars note that delegating life-and-death decisions to algorithms challenges our traditional accountability models (Source: www.mdpi.com). Many fear the moral black box: "if an AI must swerve and cause a fatality in one person's path to save five on another, how is that decision made? Who is to blame?" This so-called "moral hazard" of autonomy is unresolved legally and ethically.

In practice, any move toward autonomous cockpits will likely begin with humans closely supervising. As one NASA suggestion, we might see "human-on-the-loop" architectures, where AI performs tasks but requires periodic pilot confirmation (Source: www.mdpi.com). Early adopters might involve situations with minimal passenger risk: cargo carriers, flight testing, or remote piloting (piailots in ground stations, analogous to drone ops). Indeed, both pilots and passengers in the MDPI survey emphasized **hybrid models** first (Source: www.mdpi.com) (Source: www.mdpi.com). This might mean single-pilot desks with robust AI assistants, or remote ground-based operators that monitor a fleet of AI pilots. By scaling trust slowly, the industry hopes to avoid sudden public backlash.

Regulatory, Certification, and Policy Issues

Aviation is one of the most tightly regulated industries, and **no existing rulebook fully covers AI pilots**. Both the U.S. Federal Aviation Administration (FAA) and Europe's EASA recognize autonomy as the future, but their frameworks are at best **incipient**. As MDPI reports, *neither FAA nor EASA has a complete certification basis* for passenger-carrying autonomous aircraft (Source: www.mdpi.com). Instead, regulators are developing roadmaps: FAA's 2024 AI Safety Assurance Roadmap calls for a "progressive evidence-based certification pathway" for learning systems (Source: www.mdpi.com), and EASA's 2023 AI Concept Paper proposes a taxonomy of autonomy levels (akin to the self-driving car levels) (Source: www.mdpi.com). But translating these ideas into actual rules is just beginning.

One issue is **international harmonization**. Airlines fly globally; yet today, no two regulators agree on how to certify AI pilots, and some have not even broached the question. In fact, interviews revealed that currently **"governments have no process in place for permitting automation such as ATTOL and IAS aboard airliners."** (Source: aerospaceamerica.aiaa.org). The pilots surveyed by MDPI echoed this – they demand *"global regulatory certification"* as a prerequisite to flying pilotless (Source: www.mdpi.com). The lack of harmonized standards clearly slows investment: how can manufacturers commit to one highly local approach when others remain undefined? On the positive side, international forums are active: ICAO's Advanced Aviation Mobility panel (2024) has started drafting high-level provisions for things like "remote pilot licensing" and AI assurance (Source: www.mdpi.com), signaling that multilateral policy might eventually catch up.

Certification itself poses unique hurdles. Deterministic flight computers are usually certified once and then boxed in. For AI, a continuous-learning system could, in theory, change its decision logic after certification – a no-go in safety-critical apps. Proposed solutions include **frozen-weight models** (where the AI's neural network parameters are fixed at the point of certification) or **iterative re-certification** (airworthiness checks for the software on a regular schedule) (Source: www.mdpi.com). Both approaches are untested in cert authority practice. In addition, the FAA would likely demand vast demonstration flights akin to a driver's test for ships or cars. Sanjiv Singh (Near Earth Autonomy) envisioned an AI flight computer needing to fly "some number of kilometers and perform certain standard maneuvers" to prove reliability (Source: aerospaceamerica.aiaa.org) (Source: aerospaceamerica.aiaa.org). These tests are expensive and time-consuming, as even a modern airplane certification (with incremental improvements) can cost hundreds of millions and several years (Source: spectrum.ieee.org). One estimate noted that adding innovative autonomy features is "sure to get longer and more expensive, with no guarantee of success" (Source: spectrum.ieee.org).

Labor and insurance considerations also intersect with regulation. Pilot unions fiercely oppose crew reductions, and they hold regulatory influence. Any regulator contemplating fewer pilots will face political pressure as well as jobs concerns (Source: aerospaceamerica.aiaa.org) (Source: www.mdpi.com). Insurance firms will demand quantitative risk models; without decades of data, underwriting autonomous flights is problematic. Liability is another policy quagmire: if an AI pilot makes a bad decision, is the manufacturer at fault? The airline? The regulator for approving it? Academia suggests a "joint accountability" model, spreading liability among stakeholders based on each party's role (Source: www.mdpi.com). Some frameworks (like the EU AI Act) begin to impose requirements for transparency and documentation that could help (e.g. mandating audit trails for AI decisions) (Source: www.mdpi.com). Yet these initiatives are in early stages.

Overall, current regulatory thinking favors a **phased approach**: certify autonomy in low-risk domains first (cargo, limited routes, constrained airspace), then gradually expand to passenger service (Source: www.mdpi.com) (Source: www.mdpi.com). This matches the empirical preference of both passengers and pilots for incremental adoption (Source: www.mdpi.com) (Source: www.mdpi.com). Until a robust, standardized certification toolkit for AI exists, though, true pilotless passenger operations must remain parked on the drawing board.

Case Studies and Real-World Episodes

Examining specific incidents and experiments provides insight into why AI pilots can't simply be switched on today.

- **Asiana Airlines Flight 214 (2013)** – *Human-Autopilot Interface Failure*: A Boeing 777 on visual approach at SFO crashed into a seawall, killing 3. Investigators found the crew had mismanaged automation modes: the pilot accidentally set an autotthrottle mode incorrectly and then turned the autopilot off to hand-fly, expecting the autothrottle to maintain speed (Source: aerospaceamerica.aiaa.org). In fact, the autothrottle was off, the speed fell into a stall, and the aircraft descended into terrain. The NTSB chairman summarized: "the flight crew over-relied on automated systems they did not fully understand" (Source: aerospaceamerica.aiaa.org). This tragedy underscores that even advanced autopilot systems have complex modes which can confuse pilots. It also suggests that carelessly removing an expert human (whose *job it was to understand the automation*) could be extremely dangerous.

- **Air France Flight 447 (2009)** – *Air Data Loss and Stall*: Though not directly about the absence of pilots, AF447 is instructive. A modern Airbus A330 flying over the Atlantic encountered pitot-tube icing that caused unreliable airspeed readings. The autopilot disengaged and the crew, startled by contradictory instruments, failed to maintain proper pitch, leading to a high-altitude stall. Had an AI pilot been at the controls, would it have recognized the data failure and applied the correct recovery? We do not know, but AF447 highlights how quickly an aircraft can become

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

unflyable if flight automation goes offline unexpectedly. It emphasizes the need for any autonomous system to not only fly well *with* good data, but also to **handle misleading sensor information gracefully** – something even seasoned pilots can struggle with (Source: aerospaceamerica.aiaa.org) (Source: www.mdpi.com).

- **Qantas Flight 32 (2010)** – *Systems Reversion Challenges*: Shortly after takeoff from Singapore, an A380 suffered an uncontained engine failure that blew off part of the left wing and triggered hundreds of system warnings. The crew managed, despite the complexity, to bring the damaged plane back safely. Notably, much of the A380's automation (including its autoland) was taken offline by this event. The crew had to hand-fly much of the remainder of the flight, interpreting conflicting electronic cues (Source: aerospaceamerica.aiaa.org). For an AI pilot, this incident would be a worst-case scenario: a sudden catastrophic failure leaving partial systems online and erroneous data flooding the cockpit. Qantas 32 demonstrates the necessity of human improvisation and prioritization. An autonomous system would need robust fallback logic to detect and isolate failures without human intuition. The fact that manual piloting saved the day suggests why airlines insist on a human "pilot-in-command" until autonomy is absolutely bulletproof.

- **Airbus A350 Autonomous Taxi/Takeoff (2019)** – *Successful Pilotless Trials*: On the positive side, Airbus's ATTOL flights conclusively showed that a large modern airliner *can* taxi, take off, and land without human control under test conditions. In Toulouse, nine flights of a modified A350-1000 (with a trained safety pilot on board) used only onboard cameras and computers to identify taxiway markers and runway lines. The AI performed the entire ground and departure process, removing humans from decision-making on those phases (Source: aerospaceamerica.aiaa.org). Importantly, this experiment **did** match FAA/EASA safety procedures (by keeping a pilot ready to intervene), but it indicates that with sufficient mapping and reliable computer vision, the *routine* aspects of flight can be automated. The shortfall: this was done in good weather at one airport with a modified aircraft. The AI system was *narrowly scoped* – it only had to find painted taxi lines. For such an AI to generalize globally (night operations, snow-covered runways, atypical airport layouts) would need much more development. Still, Airbus's result is often cited as technical proof-of-concept that autonomy can extend beyond cruise. It emphasizes how far the technology has come, even if regulatory acceptance trails.

- **Company Prototypes**: A few startups have taken experimental flights. For example, Reliable Robotics (recently FAA-certified for cargo) has begun flying Cessna Caravans with an "always-on" automated pilot (Source: spectrum.ieee.org). Initially, a human pilot will remain onboard, but the system is being certified to handle most flight phases independently (including taxi, takeoff, and landing). This mirrors the aforementioned eVTOL programs: start small (general-aviation cargo, tiny eVTOLs), prove autonomy in controlled settings, then scale up. These prototypes are building the operational and credentialing experience needed for larger aircraft, but full passenger-carrying certification is still several steps away.

These cases illustrate a dual message. On one hand, **automation has scored successes**: autopilots routinely handle up to 90% of flight and can even outlast most pilots in simple tasks (Source: www.wired.com). On the other, there remain crucial junctures where the human is irreplaceable. Every major incident to date with automated components has involved either mode confusion or system failures requiring human correction. Until AI pilots can cover those edge cases as reliably as humans, it is wise to proceed cautiously.

Data Analysis and Evidence-Based Discussion

Safety and Reliability

Aviation culture demands that any new technology be demonstrated safe *beyond any doubt*. Survey data underscores that safety perception drives acceptance. In the MDPI study's passenger sample, 72% cited safety concerns about pilotless flights (Source: www.mdpi.com); no amount of hype over efficiency can outweight that fear. Pilots likewise see safety as paramount: 93% agreed that automation (like current autopilots) **improves** safety overall, but 80% *opposed* the idea of fully removing humans (Source: www.mdpi.com) (Source: www.mdpi.com). This reveals a nuanced stance: professionals welcome incremental automation benefits but firmly reject autonomy that eliminates human oversight.

Empirical analyses reflect this cautious stance. For instance, NASA's ICAROUS proved that its collision-avoidance algorithms work very well *within predefined flight domains*, but admitted that "operational reliability at airline-level safety targets ($\approx 10^{-9}$ failures per flight hour) remains unproven" (Source: www.mdpi.com). The MDPI authors synthesize this reality: **"In short, technical feasibility exists, yet operational reliability at airline-level safety targets remains unproven."** (Source: www.mdpi.com). The implication is that, even if an AI pilot could fly perfectly in 10,000 test runs, it must sustain that perfection over millions of flights in the wild. That requires *extreme* redundancy and validation. The report notes that "achieving parity with human reliability requires ultra-redundant architectures, multiple independent sensor channels, high-assurance flight control logic, and deterministic fail-safe behaviors" (Source: www.mdpi.com). These are engineering nightmares in their own right, and they illustrate why regulators are "moving slowly" – incremental proof is demanded before any full autonomy can be trusted.

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

Comparison with the automotive domain is instructive. Self-driving cars have made headlines, but even they have struggled to convincingly demonstrate 10^-9 reliability. In air travel, that threshold is non-negotiable. Therefore, experts emphasize **phased deployment**. The MDPI study and others recommend using pilotless systems first on cargo and other low-risk missions (Source: www.mdpi.com). Actual practice seems to follow: automated cargo flights (even small uncrewed freighters) are seen as the early adopters, where no human passengers means public trust is less of an issue and liability can be borne by operators. Only after hundreds of millions of safe cargo flight hours could passenger service be considered. Even then, a "hybrid" model – e.g. a remote human overseeing multiple autonomous planes – is likely to bridge the gap (Source: www.mdpi.com) (Source: www.mdpi.com).

Human-Machine Interaction

Trust is built not just on raw reliability, but on design of the human-machine interface. Studies in this report's discussion section emphasized that **both over-trust and under-trust are dangerous** (Source: www.mdpi.com). Over-trust leads to complacency (pilots not monitoring enough); under-trust means pilots would reject helpful automation entirely. Historically, aviation has learned this lesson the hard way (e.g. the Gulfstream GIII Miracle on the Hudson in 2009 involved pilots with autopilot who still had to react to unexpected geese strikes). Today's pilots voice concerns about "automation bias" and loss of situational awareness as systems become more complex (Source: www.mdpi.com). Any AI pilot must address these through clever HMI design. One concept is **adaptive automation** – systems that adjust their level of autonomy based on pilot workload and context. For example, an autopilot could be designed to occasionally ask the human to verify a decision ("Confirm pull-up command; I detect terrain and see weather to the right.") rather than retracting totally. Empirical work (outside aviation) with simulated cockpit interfaces shows that transparency – giving pilots insight into *why* the AI acted – considerably increases acceptance (Source: www.mdpi.com). This translates to calls for explainable-AI (XAI) within cockpits: if the computer could communicate its reasoning in human terms, pilots and passengers might have less anxiety.

Another factor is **human identity and employment**. Numerous pilots in surveys express not only safety anxieties, but also **fear of job loss** (Source: www.mdpi.com). This sociopsychological aspect is real. According to IATA forecasts, millions of pilot jobs will be needed for aviation growth, but the nature of those jobs may change (Source: www.mdpi.com). The MDPI interviews noted that about 40% of pilots expected *new roles* (e.g. in supervision or AI oversight) to emerge as autonomy grows (Source: www.mdpi.com). Acceptance strategies thus include emphasizing that pilots will "partner" with AI, rather than be replaced by it (Source: www.mdpi.com). Legally and ethically, there is also pressure to keep a human "in the loop" (either in cockpit or remotely) for a long time (Source: aerospaceamerica.aiaa.org). The aviation profession respects tradition and training; retraining the workforce to manage autonomous systems (e.g. as remote supervisors) will be a major policy task.

Liberties, Liability, and Ethics

In aviation today, the pilot-in-command holds ultimate legal responsibility for the flight. An AI pilot would upend that doctrine. Who is liable if "the autopilot made the right decision 99 times but decided incorrectly the 100th time"? Elbasyouny & Dababneh explain that in unmanned systems, accountability will diffuse across designers, operators, manufacturers, and regulators (Source: www.mdpi.com). Some scholars argue for **joint liability frameworks**, where legal responsibility is allocated proportionally by contributing parties (Source: www.mdpi.com). However, courts and insurance companies have not fully sorted this out. In practice, an international treaty or agreement (ratified by civil aviation authorities) might eventually clarify that a technologically advanced AI performs as the legal "pilot," making the airline or manufacturer akin to today's airlines bearing liability for pilot error. Until such doctrines are written into law, the specter of lawsuits could make carriers and manufacturers extremely cautious about deploying untested AI systems, even on cargo flights.

In summary, then, the **human dimension** is as much a roadblock as the technical one. People want absolute safety assurances and the presence of a trained human. They want transparency and accountability. Even if engineers can solve every control equation, the industry must also manage public perception, workforce transitions, and legal structures. This complex mosaic of requirements – some intangible, some technical – is a critical reason why AI pilots remain grounded for the foreseeable future.

Economic and Operational Considerations

From an economic standpoint, doing away with pilots offers **attractive cost savings**: each pilot removed saves salary, training, and benefits. Elbasyouny & Dababneh estimate airlines could eventually cut 10–15% off operating costs by eliminating crew (Source: www.mdpi.com). That could translate to billions annually for major carriers. However, these savings are *theoretical* and hinge on several conditions. Most importantly, the "safety trade-off" is the non-negotiable bottleneck for passengers and regulators (Source: www.mdpi.com) – in surveys both groups ranked safety far above

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

cost when considering pilotless flights. Additionally, significant **infrastructure investments** are needed to realize even partial autonomy (Source: www.mdpi.com): high-bandwidth datalinks to maintain communication with ground stations, robust cybersecurity grids, satellite uplinks for navigation redundancy, plus new sorts of ground-based air traffic control integration. These all incur capital and operational expenditures that offset crew savings.

Cybersecurity is again a concern here. Studies warn that the costs of shoring up connectivity may exceed the benefits: if every autonomous plane must have deep packet inspection, encryption, and intrusion prevention systems, the extra maintenance and retrofitting can negate the crew-cost savings (Source: www.mdpi.com). Similarly, insurers may demand higher premiums (at least initially) for autonomous operations until they gather long-term data. On the other hand, some operational advantages of autonomy could improve economics: autonomous planes can fly longer segments without pilot duty-time restrictions, enabling more efficient scheduling and usage of aircraft (Source: www.mdpi.com). They could also potentially operate at odd hours or on under-served routes where pilot supply is tight.

Smart business strategies recognize that *full* autonomy is far out, but *partial* autonomy can already deliver ROI. Many airlines and OEMs are exploring **single-pilot cockpits** with a ground support pilot (essentially remote co-pilot). For example, some regional airlines trialed flying a two-pilot plane with only one pilot physically onboard and one monitoring from headquarters via video/audio link (still within current legal confines). If a remote co-pilot can handle tasks digitally, the true pilot can have more rest periods. Meanwhile, manufacturers have highlighted how features like Airbus's ATTOL program (autonomous taxi/takeoff/landing) can yield efficiency improvements without eliminating humans (Source: www.mdpi.com). In the water, automation can optimize flight profiles and energy usage (e.g. continuous descent approaches, more precise speed schedules). These incremental gains improve fuel efficiency and maintenance scheduling.

Table 3 illustrates the **operational logic** of gradual autonomy deployment (simplified):

PHASE	OPERATION TYPE	ANALYSIS & AIRCREW
1. Remote/AI-augmented Cargo	Uncrewed cargo planes on fixed routes, controlled by remote pilots in a ground station.	Low passenger risk, tests datalinks, initial trust-building. Commercial cargo shipments only.
2. Single-Pilot Commercial	Passenger flights with one pilot in cockpit and remote co-pilot.	Pilot's workload balanced by autopilot; legal pilot in cabin. Reduces cockpit staffing costs moderately.
3. Ground-Supervised Commuter	Small passenger aircraft (seats 2–6) piloted remotely under visual line of sight or within narrow corridors.	Possibly no one on board. Testbed for urban air mobility eVTOLs, air taxis, or shuttles.
4. Fully Unpiloted Commercial (Long Term)	Widebody/jets flying major routes with AI pilots, no humans aboard.	Depends on decades of safe history, fully adapted regs, regained public trust.

Each progressive step involves accumulating safety data and public familiarity. Today's mindset in industry is that Stage 1 (cargo drones/robotics) is underway, Stage 2 (single-pilot, heavy automation with oversight) may come next, and Stage 4 (nightmare scenario of no pilots on 777s) is decades away. Economic incentives exist at every stage, but they are tightly coupled with technological readiness and regulatory permission.

Future Directions and Conclusions

Where does all this lead? Most analysts agree that we will not see pilotless passenger jets in the next few decades – perhaps not before mid-century. Europe's EASA explicitly expects that fully autonomous airliners are unlikely until after 2050 (Source: spectrum.ieee.org), barring radical breakthroughs. Technology continues to advance – AI is improving, quantum computing and new sensors are on the horizon – but the constraints are not only technical. The consensus, echoed by regulators and experts, is for **evolution, not revolution**. Autonomous capability will enter the cockpit piecemeal: better vision systems (as Airbus did), improved decision-support tools, more resilient flight controls, and advanced "digital copilots." Human pilots will remain central, albeit with shifting roles toward system management, emergency piloting when needed, and oversight of automation.

Meanwhile, focused deployments (cargo, short-range eVTOL, and remote operations) will serve as proving grounds. Flight data recorders from these operations will accumulate evidence of safety. Regulators will watch closely; every software update, every anomaly, will be evidence studied for years. New certification standards for AI are likely to emerge: we may see the FAA adopt "machine learning software assurance" guidelines, or EASA integrate autonomy levels into certification classes. International bodies like ICAO could set global metrics for AI-pilot reliability.

LANDING.AERO

Why Can't AI Fly a Plane? Technical & Human Barriers

From a social perspective, public opinion must shift. Some experts suggest framing AI pilots as continuations of the trend toward automation, comparing them to autopilots that have always flown much of the trip. Transparency campaigns might emphasize that human judgment remains vital (even if offsite), to ease the perception of risk (Source: www.mdpi.com). Trust cannot be rushed; it will need time as successes accumulate under careful, visible guidance.

Data and systems integration will need advancement too. For AI to function in the national airspace, flights will likely be highly networked (e.g. Automatic Dependent Surveillance–Broadcast, space-based ADS-B for precise tracking, and common avionics datalinks). If ground control centers will manage fleets of autonomous planes, these centers will become as rigorous as existing control towers, with stringent messaging and backup systems.

Finally, it is worth stressing that autonomous flight *may not eliminate humans* so much as transform them. The MDPI study concludes that autonomy will likely "redefine" the human role, not eradicate it (Source: www.mdpi.com). Just as the first autopilots shifted pilots from hand-flying to supervision, AI pilots will require humans to become supervisors of supervisors (pilots of virtual pilots?). The "uncanny valley" analogy applies: partial autonomy can actually undermine confidence if not accompanied by transparency. The only way out, as MDPI suggests, is to "keep going" in small steps – to build the AI-brain carefully, but always with a human tandem, until the technology finally lands where the people feel safe to get off-board.

In conclusion, current barriers to AI piloting aircraft are multi-dimensional. Technical maturity is improving, but the gap to "airline-grade" autonomy remains. Certification hurdles are looming and unresolved. Simultaneously, societal factors – trust, ethics, regulation – impose their own limits. Our analysis shows that **"what's stopping AI from flying a plane"** is not a single obstacle but a complex chain: data and machine learning limits, safety and reliability standards, human factors and trust, regulatory and liability frameworks, and cost/benefit calculations. Each link in this chain must be addressed. For now, the safest path forward is a cautious hybrid approach where AI supplements rather than supplants human pilots, allowing industry and public alike to adapt while ensuring that the safety of flight remains our highest priority (Source: www.mdpi.com) (Source: www.mdpi.com).

Future Work: Continued research should focus on specific enabling technologies (robust sensor fusion, explainability, certification methodologies) and longitudinal studies of trust and behavior in mixed-autonomy cockpits. Regulatory agencies will also need to iterate on concrete standards for "learning-enabled flight" and evaluate them through simulations and test flights. Only by attacking this problem systematically – and empirically – can humanity turn the question on its head and make autonomous flight *truly* as safe (or safer) than having pilots at the helm.

References: All claims above are supported by industry reports, academic studies, and expert commentary. Notable sources include investigation reports (NTSB on Asiana 214), aerospace industry journalism (Source: aerospaceamerica.aiaa.org) (Source: aerospaceamerica.aiaa.org), FAA/EASA planning documents (Source: www.mdpi.com) (Source: www.mdpi.com), and peer-reviewed analysis of autonomy in aviation (Source: www.mdpi.com) (Source: www.mdpi.com), among others cited inline. These outline the current state of technology, documented incidents, survey statistics, and projected roadmaps, which underlie the conclusions here.

Tags: ai in aviation, autonomous flight, ai pilot, pilotless aircraft, flight automation, aviation safety, aviation regulation, do-178c

---