

Boeing, Airbus and GUDIYA - Control Philosophy

Who has the final say at the edge of the envelope

A comparison of Boeing's conventional-looking control column and Airbus's sidestick raises a question much larger than cockpit ergonomics: when an operator commands a maneuver near the boundary of safe operation, should the system *warn and preserve the command*, or *interpret and constrain the command*?

That question is central to GUDIYA. An autonomous actant, like a pilot, has a destination and a reason for choosing its next move. But the actant may not see every other participant, coupling, constraint, or emerging hazard in the Cognitive Field. It can be competent at pursuing an objective while being poorly oriented to the stability of the larger system.

The useful lesson is not that Boeing is unsafe or that Airbus removes the pilot. Both manufacturers use sophisticated protections, and the details vary by aircraft and control mode. The contrast is one of *emphasis*: preserve operator agency and make boundaries perceptible, or mediate commands so that normal operation remains inside a protected envelope. In the environment of HCAS (dense and machine-speed operating actants environment) GUDIYA grid needs the latter as a governing principle—and the former as an indispensable source of legibility.



A careful reading of the aircraft analogy

On many Boeing aircraft, the paired control columns move together. Their motion provides the crew with a visible and tactile indication of the other's input and, depending on the aircraft, of automated control activity. A column is not, however, a direct mechanical window into the aerodynamic loads on every modern Boeing. The 777 and 787 employ fly-by-wire systems; the pilot is not literally holding the control surfaces through uninterrupted cables.

Airbus fly-by-wire aircraft commonly use sidesticks that are not mechanically linked to each other. In normal law, a pilot's input is interpreted by flight-control computers, which implement protections against certain flight-envelope exceedances. The absence of linked sticks creates a human-factors burden of its own: one pilot cannot infer the other pilot's input merely by looking at or feeling their own stick. Airbus addresses simultaneous inputs with alerts and priority controls.

Nor is the Airbus envelope an unconditional wall. Flight-control degradations can change the applicable law and remove some protections. An envelope-protection philosophy therefore depends on sensing, redundancy, annunciation, crew understanding, and procedures for degraded modes.

The underlying question is not **column versus joystick**. It is:

How does a control architecture allocate initiative, information, authority, and responsibility between the operator and the stabilizing system?

The GUDIYA answer

GUDIYA should not choose every action for every actant. That would sacrifice the adaptive intelligence that makes agentic systems valuable. An actant should normally be free to plan, explore, and act within its authorized objective and operating envelope.

But permission to request an action is not permission to destabilize the Cognitive Field. When a proposed action or coupling approaches a boundary, the Grid must be able to reshape, delay, rate-limit, stage, redirect, quarantine, or deny it. The stabilizing layer is not merely a warning light mounted next to the actant; it is part of the effective control path.

In shorthand:

Actant intent → Grid envelope check → permitted or modified execution → observed outcome

This is why GUDIYA is more Airbus-like. The actant remains the source of initiative, but its command is mediated by a control law that preserves a higher-order stability constraint.

There is a crucial difference, however. An aircraft's certified flight envelope concerns a comparatively well-characterized physical vehicle. GUDIYA may need to constrain an actant on behalf of an interaction, an enterprise, a fleet, a market, or an entire national Cognitive Field. Its boundaries may be changing, partly inferred, contested among stakeholders, and different at different layers of Stability Authority. Some limits will be hard constraints; others will be probabilistic and require uncertainty margins.

GUDIYA therefore cannot simply copy a cockpit. It must engineer a field-referenced cognitive envelope.

The Stability Envelope Requires an External Reference

An actant can know its objective, permissions, tool limits, and internal state. It can even carry fixed rules, or guardrails, that prohibit particular actions. But these facts do not tell it where the boundary of system-wide stability lies. That boundary also depends on other actants, shared resources, changing conditions, and the effects of their interactions. It may also depend on which layer of the AI Economy the actant is operating in: the Agentic Layer, the Digital Services Layer, or the Resource Layer. (<https://www.manhattanproject20.com/blogs/the-ai-economy-architecture-v3>)

Two actants may each operate within their individual limits while converging on an unstable rendezvous. A sequence of individually permissible actions may create a dangerous collective trajectory. And an action that appears acceptable within the Agentic Layer may acquire different consequences when it invokes a digital service or commits scarce physical resources. The stability envelope must account for those cross-layer effects as well as the actant's immediate behavior.

GUDIYA therefore needs an external reference for two purposes. First, the Grid supplies orientation: it observes the actant's motion relative to other actants, AI Economy layers, and the wider Cognitive Field. Second, it supplies authority: it determines which system and stakeholder constraints govern that motion. Neither function can be reliably derived from the actant's self-assessment alone.

An actant can carry local guardrails, but it cannot independently determine a dynamically valid, field-wide stability envelope. That envelope must be externally referenced and continuously assessed across actants, interactions, and AI Economy layers.

The missing piece

Envelope protection creates its own hazard if the operator ceases to perceive the boundary. An actant that is silently corrected every time it approaches instability may continue planning as though its commands were executed exactly as intended. A human supervisor may believe the operation is comfortably stable when it has repeatedly required Grid rescue.

GUDIYA should preserve the valuable part of the Boeing intuition: **the operator must feel what the system is doing**—not necessarily through a physical control column, but through intelligible telemetry.

An actant and its authorized human supervisors should be able to learn, as appropriate:

- which envelope boundary is approaching;
- the estimated margin and uncertainty;
- which command was modified and why;
- whether another actant's input or a collective trajectory caused the constraint;
- whether Grid protection is healthy, degraded, or unavailable;
- what safe alternatives remain; and
- whether the intervention achieved a verified stable state.

The Cognitive Attitude Indicator (CAI) (<https://www.manhattanproject20.com/blogs/cognitive-attitude-indicator>) can serve as the actant's artificial horizon. GCR's intervention record can serve as an auditable account of what was demanded, what was permitted, and what happened. The GESR can retain this history across model, host, and CEE changes. This is **cognitive control feel**: a faithful

representation of external stability forces and margins, even when there is no physical force on a stick.

The design objective is not to make every actant negotiate every minor correction. It is to prevent silent, habitual dependence on invisible rescue.

Authority not paternalism

An envelope has meaning only if its authority is defined. GUDIYA cannot let an actant redefine a critical boundary merely because crossing it would help finish its task. Conversely, a Grid operator must not turn the word *stability* into an unaccountable veto over legitimate activity.

Control authority should therefore be tiered and reviewable:

- | | |
|---|---|
| 1. Normal operation: | the actant retains initiative inside its authorized envelope. |
| 2. Boundary approach: | the Grid supplies orientation and may shape or pace the maneuver. |
| 3. Predicted exceedance: | the Grid constrains or refuses the action under an explicit control rule. |
| 4. Degraded sensing or protection: | the actant enters a narrower fallback envelope, seeks a safe harbor, or pauses high-consequence coupling. |
| 5. Emergency: | GCR invokes proportionate arrestment, quarantine, or SCRAM, with independent logging and subsequent review. |

The strongest boundaries must derive from an accountable hierarchy of authority—not from a hidden preference of the model developer, actant owner, or Grid operator alone. GUDIYA's proposed Cognitive Utilities Commission (<https://www.manhattanproject20.com/blogs/cognitive-utilities-commission>) matters precisely because protection of the 99% cannot rest entirely with those who benefit from moving fastest.

A control system must be honest about what it cannot see

Airbus-style envelope protection is attractive only if the system's estimate of the envelope can be trusted. Faulty sensors, contradictory data, delayed telemetry, unknown actants, and rapidly emerging couplings can make a confident intervention wrong—or cause an intervention to arrive too late.

The Grid must therefore distinguish **known-safe**, **known-unsafe**, and **insufficiently observed** conditions. Uncertainty is itself a control input. When it grows, GUDIYA should not claim omniscience; it should reduce permitted autonomy, shorten observation intervals, demand a rendezvous with an authoritative reference, or divert the actant to a Cognitive Safe Harbor.

This also connects to Agentive ETOPS (metaphor for FAA's regulation of twinjet airplanes). The actant's operating clearance should depend on how far it can move through cognitive space while retaining a reachable stabilization or recovery point. A long-running actant does not earn unlimited clearance merely because it has operated for many hours without incident.

The control philosophy in one sentence

GUDIYA should permit ambitious actant initiative within a field-referenced stability envelope, constrain commands that would cross that envelope, and make the boundary, the intervention, and the remaining margin perceptible and accountable.

The aircraft analogy now yields a synthesis rather than a brand preference:

- Airbus contributes the principle of active envelope governance.
- Boeing's linked-control intuition contributes the demand for visible, comprehensible control action.
- GUDIYA needs both—because an invisible stabilizer can become a source of dependency, while a visible warning without effective authority may arrive too late.

That is the transition from merely supervising an intelligent actant to engineering a stable Cognitive Field around it.

Sources and analogy note

- Airbus, “Flight Envelope Protection”, on normal-law limits and degraded protections.
- Airbus, “Fly-by-wire”, including Boeing 777/787 fly-by-wire systems.
- Airbus Flight Safety, “What about Dual Sidestick Input?”, on input priority and annunciation.
- Boeing, 737 MAX return-to-service information, illustrating its explicit attention to pilot override in that aircraft's revised augmentation system.

The GUDIYA architecture and terminology in this chapter are proposed design principles, not claims made by the aircraft manufacturers.