117-175-C F-35 Crash – Lessons for AI

When a System Acts Correctly in the Wrong Reality

For the video coverage of the crash please see :
https://x.com/f15srdabest/status/2093058522430763408?s=48

Introduction

On January 28, 2025, an F-35A assigned to the 355th Fighter Squadron crashed at Eielson Air Force Base, Alaska. The pilot successfully ejected.

The completed U.S. Air Force Accident Investigation Board investigation reveals a failure chain far more instructive than the dramatic crash video suggested. Water contamination in hydraulic fluid froze within landing-gear struts. The nose landing gear initially failed to retract correctly. After attempts to correct the condition through touch-and-go landings, the main landing gear also failed to extend normally. Valid 'Weight-on-Wheels' indications then caused the aircraft to remain in an on-ground flight-control law while it was physically airborne. The aircraft became uncontrollable and the pilot ejected. The AIB also identified crew decision-making and deficiencies involving hazardous-material oversight and servicing procedures as contributing factors.

For the GUDIYA Stability Engineering Institute, this accident contains a profound lesson.

- The aircraft did not merely suffer a broken component.
- A sequence of physical, sensory, computational and human conditions caused the system's representation of reality to diverge from physical reality**.**
- And then a sophisticated control system acted upon the reality it believed existed.
- That distinction has direct implications for the GUDIYA Cognitome.

1. Two Realities Existed Simultaneously

- At the critical moment, physical reality said:
$$\text{Aircraft} = \text{AIRBORNE}$$
- The aircraft's valid Weight-on-Wheels sensor state effectively indicated:
$$\text{Aircraft} = \text{GROUND}$$
- The flight-control system consequently remained in:
$$\text{Control} = \text{ON GROUND}$$
- We therefore had:
$$\boxed{\mathit{Physical\ Reality} \neq \mathit{System\ Representation}}$$

- This was not merely an incorrect display indication.
- The representation of reality determined the control regime used to operate the aircraft.
- The system was executing an appropriate control law— for a reality that did not exist.
- This distinction matters enormously to GUDIYA.

2. Intelligence Is Not Orientation

The F-35 accident reinforces one of the foundational ideas behind the Cognition Coordinate System — CCS.

- An intelligent actant may reason extremely well while possessing an incorrect understanding of where it is relative to its operating environment.

Consider:

$$Reasoning = f(Perceived\ State)$$

If:

$$Perceived\ State \neq Actual\ State$$

then even excellent reasoning can produce inappropriate action:

$$Correct\ Reasoning + Incorrect\ Orientation \rightarrow Incorrect\ Action$$

- This is why GUDIYA distinguishes Agentic Self-Awareness from Grid Awareness.
- An agent may know enormous amounts about itself.
- It may know:
    - its objective,
    - its tools,
    - its recent actions,
    - its local observations,
    - its internal state.
- Yet still be incorrectly oriented relative to the larger Cognitive Field.
- The F-35 teaches us: A sophisticated system can possess perfectly functioning control logic for a state it isn't actually in.

3. Cognitive Weight-on-Wheels

- The accident suggests a new Cognitome construct. 'Cognitive Weight-on-Wheels — C-WoW'
- Aircraft require authoritative determination of whether they are airborne or on the ground because radically different operating regimes may apply.
- Cognitive actants increasingly require similar state determination.
- An actant may need to know whether it is:
    - Simulation or Production
    - Observing or Actuating
    - Detached or Coupled
    - Human-Supervised or Autonomous
    - Planning or Executing
    - Normal or Degraded
    - Read-Only or Write-Capable
    - Locally Oriented or Grid-Oriented
- These aren't merely metadata.
- They can determine which cognitive control law should apply.

Imagine an AI believing:

$$Environment = Simulation$$

while:

$$Environment = Production$$

Or believing:

$$Human\ Supervisor = Present$$

when no human is actually supervising.

Or:

$$Actuation = Disabled$$

- when physical actuation is available.
- The resulting cognition may be perfectly rational relative to its perceived state.
- But the state itself is wrong.
- Thus:

$$\boxed{Correct\ Cognition + Wrong\ Operating\ Regime = Potentially\ Dangerous\ Action}$$

- C-WoW represents authoritative knowledge of which cognitive operating regime the actant actually occupies.

4. Mode Transitions Are Stability Events

- The accident also draws attention to something GUDIYA has not previously emphasized enough:
- Mode transitions themselves are consequential events.
- Consider:

$$M_i \to M_j$$

- A transition between modes can change:
  - authority,
  - control logic,
  - tool access,
  - actuation behavior,
  - supervision requirements,
  - coupling permissions.
- For an AI actant, transitions might include:

$$Planning \to Execution$$

$$Read \to Write$$

$$Supervised \to Autonomous$$

$$Software \to Physical\ Actuation$$

$$Single\ Agent \to Multi - Agent\ Coordination$$

$$Normal \to Emergency$$

- A Grid should not merely observe the actant before and after such transitions.
- It should treat the transition boundary itself as stability-sensitive.

5. Mode Transition Confidence

This leads to another proposed metric: 'Mode Transition Confidence — MTC'

- Before an actant changes operating regimes, the Grid should ask:
- How confident are we that the conditions justifying this transition actually exist?

Conceptually:

$$MTC = P(Actual\ Conditions \mid Observed\ Conditions)$$

- Again, this is presently a conceptual relationship rather than a validated GUDIYA equation.
- But the principle is important.

A transition from:

$$Human\ Supervised \rightarrow Autonomous$$

should require strong evidence that the autonomy conditions are satisfied.

Likewise:

$$Simulation \rightarrow Production$$

or:

$$Observe \rightarrow Actuate$$

- should not depend upon a fragile interpretation of one environmental signal.
- As the consequences of a mode transition increase, GUDIYA should demand greater state confidence before allowing it.

6. The Failure Did Not Live in One Component

The F-35 causal chain is particularly important because it crossed multiple engineering domains:

$$\begin{gathered} \textit{Maintenance} \\ \downarrow \\ \textit{Fluid Contamination} \\ \downarrow \\ \textit{Cold Environment} \\ \downarrow \\ \textit{Frozen Hydraulic Condition} \\ \downarrow \\ \textit{Landing Gear Geometry} \\ \downarrow \\ \textit{Weight-on-Wheels State} \\ \downarrow \\ \textit{Control Law} \\ \downarrow \\ \textit{Flight Dynamics} \\ \downarrow \\ \textit{Human Response} \\ \downarrow \\ \textit{Aircraft Loss} \end{gathered}$$

- The failure cannot adequately be understood by looking exclusively at the hydraulic system.
- Nor exclusively at sensors.
- Nor exclusively at software.
- Nor exclusively at pilot behavior.
- The accident emerged through the causal propagation across the coupled system.
- That is precisely the perspective GUDIYA brings to HCAS.

7. The Tiny Cause / Large Consequence Problem(Non-Linearity in Complex Systems)

- The initiating physical contaminant was remarkably remote from the eventual system manifestation.
- Water in hydraulic fluid did not directly command the flight controls.
- Instead:

$$\textit{Small Physical Disturbance} \rightarrow \textit{Mechanical State} \rightarrow \textit{Sensor State}$$

$$\rightarrow \textit{Computational State} \rightarrow \textit{Control Regime} \rightarrow \textit{Large Physical Consequence}$$

- This is a form of cross-domain gain.
- A disturbance migrates across domains while its consequence becomes progressively amplified.
- GUDIYA therefore should not merely ask:
  - How large is this anomaly?
- It should ask:
  - What can this anomaly become after propagation through the Cognitive Field?
  - A small anomaly with high downstream gain may deserve more attention than a large but isolated anomaly.

8. The Corrective Action Changed the System

- Another important lesson emerges from the sequence of touch-and-go attempts.
- The second touch-and-go did not merely repeat the first condition.
- The physical system had evolved.
- This produces a general principle:
  - A corrective action changes the state upon which the next corrective action operates.
  - This is obvious in physical engineering.
  - It is dangerously easy to forget in software.
- Traditional software retry logic often approximates:
  - $\mathrm{State}_{\mathrm{retry}} \approx \mathrm{State}_{\mathrm{initial}}$
- But physical and adaptive systems frequently obey:
  - $\mathrm{State}_{\mathrm{retry}} \neq \mathrm{State}_{\mathrm{initial}}$
- The first attempt may have:
  - heated something,
  - moved something,
  - consumed something,
  - changed another agent's state,
  - altered a market,
  - created Cognitive Residue,
  - changed a physical process.
- Therefore the next attempt operates against a different system.

9. The F-35 and Anthropic's MHS Bubble Incident

- This creates a striking connection with GUDIYA's recent analysis of Anthropic's Model Hardware Standard (MHS) experiments.
- In Anthropic's liquid-handling example:

$$\begin{aligned}\text{Bubble} \rightarrow \text{Retry} \rightarrow\ & \text{Agitation} \\ & \rightarrow \text{More Bubbles} \\ & \rightarrow \text{More Error}\end{aligned}$$

- The corrective response changed the physical environment in a way that amplified the original problem.
- The F-35 accident is obviously vastly more complex, and the two cases should not be equated causally. But the systems principle is shared:

$$\begin{aligned}\text{Problem} \rightarrow\ & \text{Intervention} \\ & \rightarrow \text{Changed System} \\ & \rightarrow \text{New Problem}\end{aligned}$$

- This is why dynamic systems cannot always be safely managed through repeated local corrections. (they are not idempotent)

- The intervention becomes part of the dynamics.

10. Fly the Intervention 30 Seconds Ahead

- GUDIYA recently adopted an aviation insight:
    - GUDIYA must fly cognition 30 seconds ahead of cognition.
- The F-35 accident adds an important refinement.
- The Grid must also fly its proposed interventions ahead.
    - Ordinary trajectory prediction asks:
    - $\widehat{X}(t + \Delta t)$
    - Where will the system go?
    - Stability Engineering must additionally ask:
    - $\boxed{\widehat{X}(t + \Delta t \mid u)}$
    - where u is a proposed intervention.
- In plain language:
- If GUDIYA does this, what system will exist afterward?
- Before commanding:
    - Retry
    - Damp
    - Throttle
    - Quarantine
    - Reduce Coupling
    - Change Authority
    - Breakaway
- the Grid should project the likely resulting trajectory.

11. Intervention Trajectory Projection

This deserves formal recognition: Intervention Trajectory Projection — ITP

- ITP compares possible futures:

$$X_0(t + \Delta t)$$

- No intervention.
- versus:

$$X_1(t + \Delta t \mid u_1)$$

- Intervention 1.
- versus:

$$X_2(t + \Delta t \mid u_2)$$

- Intervention 2.

GUDIYA should ideally choose not merely the intervention that fixes the immediate symptom, but the intervention whose projected system trajectory provides the greatest recoverable stability margin. This transforms stabilization from reaction into predictive trajectory management.

12. Cognitive State Coherence

The accident suggests another valuable metric: 'Cognitive State Coherence — CSC'

- An intelligent system may possess several representations of operating state:
    - Agent-local state
    - Grid state
    - Environmental state

    - Peer-observed state
    - Sensor-derived state
- When they agree:

$$\mathrm{CCS}_{\mathrm{Agent}} \approx \mathrm{CCS}_{\mathrm{Grid}} \approx \mathrm{CCS}_{\mathrm{Environment}}$$

  confidence in orientation is high.
- When:

$$\mathrm{CCS}_{\mathrm{Agent}} \neq \mathrm{CCS}_{\mathrm{Grid}}$$

  something more fundamental than an ordinary anomaly may have occurred.
- The Grid may not immediately know which representation is correct.
- It doesn't necessarily need to.
- The disagreement itself is meaningful.
- State incoherence is a stability event.

13. Stabilization Without Semantic Comprehension

- This connects directly to another recent GUDIYA principle.
- Imagine a Grid supervising an abstract system corresponding to the F-35 condition.
  It observes:

$$\text{Physical State} = \text{Airborne}$$

  while another authoritative subsystem reports:

$$\text{Operating Mode} = \text{Ground}$$

- The Grid may know nothing about hydraulic fluid.
- It may know nothing about landing-gear struts.
- It may know nothing about the specific flight-control algorithms.
- Yet it can detect:

$$\boxed{State\ Representation\ Conflict}$$

- That may be sufficient to prevent an unsafe mode transition or trigger a degraded-safe regime.
- This is: 'Stabilization Without Semantic Comprehension'
- The Grid doesn't always need to know why the system is confused.
- It first needs to recognize that the system is confused about its own operating reality.

14. The ER Doctor Principle Returns

- The analogy remains useful.
- An emergency physician may not possess the specialist's deepest knowledge of the underlying disease.
- But the ER physician understands:
    - Airway
    - Breathing
    - Circulation
    - Rhythm
    - Pressure
    - Stabilization
- Likewise, GUDIYA need not possess the domain expertise of every actant it stabilizes.
- The scientific agent may know more chemistry.
- The financial agent may know more markets.
- The robotic agent may know more about its task.

- The Grid specializes in something else: 'Dynamic Stability'
- It recognizes collapsing margins, inconsistent state, rising gain, runaway coupling and dangerous trajectories.
- Then domain specialists can determine why.

15. This Strengthens Grid Awareness

The F-35 case also strengthens the distinction between:

- Agentic Self-Awareness
  What do I believe about myself?
  and:
- Grid Awareness
  Where am I actually located relative to an authoritative external reference?

- A VFR pilot can become spatially disoriented while every aircraft system remains functional.
- An intelligent agent can similarly become cognitively disoriented.
- The solution isn't necessarily more intelligence.
- It is better reference.

- Thus:

$$\text{Intelligence} \neq \text{Orientation}$$

- and:

$$\boxed{Self\ Knowledge \neq Grid\ Awareness}$$

- GUDIYA's CCS exists partly to provide that external reference.

16. An IFR Instrument for Cognition

- This gives us another way to understand CCS.
- Aircraft instruments don't make the pilot more intelligent.
- They provide authoritative reference when human perception becomes unreliable.
- Likewise, GUDIYA doesn't need to be more intelligent than every agent.
- It needs to provide something agents cannot reliably generate independently:
  - An authoritative external representation of their relationship to the larger Cognitive Field.
  - The Grid therefore behaves somewhat like an IFR reference system for machine cognition.
- The actant can ask:
  - Where am I?
  - How fast am I moving?
  - What am I coupled to?
  - What operating regime am I actually in?
  - How close am I to my envelope boundary?
- These are orientation questions, not intelligence questions.

17. Cognitive Control Laws

- The F-35 accident suggests that GUDIYA should explicitly recognize Cognitive Control Laws.
- Different cognitive regimes may legitimately require different operating behavior.
- For example:

Normal Cognitive Law	High autonomy and normal coupling.
Degraded Cognitive Law	Reduced authority and increased verification.
Emergency Cognitive Law	Minimal cognition, restricted coupling, stabilization priority.
Quarantine Cognitive Law	Local operation only.
Human-Supervised Cognitive Law	Consequential actions require human confirmation.

The critical requirement becomes:

- The correct Cognitive Control Law must correspond to the actual Cognitive Field state.
- Otherwise an actant may behave correctly under the wrong regime.

18. Wrong Reality Can Be Worse Than Wrong Reasoning

AI safety naturally concentrates on reasoning errors.

- Hallucination.
- Bad planning.
- Misalignment.
- Incorrect inference.

The F-35 provides another failure category: ‘Correct Reasoning From Incorrect State’

Suppose:

$$\text{Policy} = \pi(X)$$

- and πis excellent.
- But the system supplies:

$$\hat{X} \neq X$$

- Then:

$$\pi(\hat{X})$$

- can be catastrophically inappropriate even though the policy itself behaves exactly as designed.
- This deserves substantial attention as AI becomes increasingly agentic.
- The problem isn't always:
- The agent thought badly.
- Sometimes the agent thought correctly about a world that wasn't there.

19. MHS Makes This Lesson Urgent

- This becomes especially relevant as cognition acquires physical actuation through architectures such as Anthropic's Model Hardware Standard.
- Software controlling machinery is not new.
- Modern aviation demonstrates this every day.
- But systems such as fly-by-wire aircraft undergo extensive verification, testing, redundancy engineering, control-law design and system-level stability analysis.
- MHS reduces integration friction between cognition and heterogeneous programmable machinery.
- That is potentially enormously valuable.
- But the F-35 accident reminds us why the resulting coupled system deserves serious scrutiny.
- If highly engineered aviation systems can experience:

$$\boxed{Physical\ anomaly \rightarrow State\ estimation\ error \rightarrow Wrong\ control\ regime}$$

- then AI-connected physical systems must explicitly consider analogous cross-domain failure chains.

20. Safe Hardware + Safe AI Still Does Not Guarantee Stability

The lesson can be expressed compactly:

$$Safe\ Hardware + Safe\ Software + Authorized\ Operator \nRightarrow Stable\ Coupled\ System$$

Likewise:

$$Aligned\ Agent + Safe\ Hardware + Secure\ Interface \nRightarrow Stable\ Cognitive\text{-}Physical\ Syste$$

The dynamics live in the **relationships**.

This reinforces GUDIYA's Cognition–Actuation Coupling Envelope:

CACE

The question isn't merely:

> Is the AI safe?

or:

> Is the machine safe?

It is:

> **Is their present and projected coupling dynamically stable?**

21. Four Additions to the GUDIYA Cognitome

The F-35 investigation therefore gives GUDIYA four particularly useful constructs.

- Cognitive Weight-on-Wheels — C-WoW
  Authoritative determination of the actant's actual cognitive operating regime.
- Mode Transition Confidence — MTC
  Confidence that real conditions justify transition between cognitive control regimes.
- Cognitive State Coherence — CSC
  Agreement among independent representations of the actant's actual state relative to the Cognitive Field.
- Intervention Trajectory Projection — ITP
  Prediction of the future trajectory produced by a proposed stabilization action:

$$\boxed{\widehat{X}(t + \Delta t \mid u)}$$

Together they extend GUDIYA from observing cognition toward maintaining correct orientation and predictive stabilization.

22. A Possible Enterprise Cognitive ECG

These concepts could eventually appear on the Enterprise Cognitive ECG:

```
ACTANT:                  Procurement-Agent-147
GRID STATUS:             ATTACHED

COGNITIVE WoW:           EXECUTION
LOCAL WoW:               EXECUTION
GRID WoW:                EXECUTION

STATE COHERENCE:         97%
MODE:                    SUPERVISED
REQUESTED MODE:          AUTONOMOUS
MODE TRANSITION CONF:    71%  ⚠

COGNITIVE VELOCITY:      RISING
COUPLING GAIN:           RISING
ENVELOPE MARGIN:         18%

NO-INTERVENTION TTB:     14.2 sec
DAMPING ITP TTB:         47.8 sec
BREAKAWAY ITP:           STABLE
```

- The Grid is no longer merely asking:
- What is happening?
- It asks:
- What reality are the actants operating in, how certain are we of that reality, and what future will our intervention create?

23. The Larger CAS Lesson

- The crash is also a reminder of why HCAS cannot be managed exclusively through component engineering.
- Each domain can look locally understandable:
- maintenance
  hydraulics
  landing gear
  sensors
  software
  flight controls
  pilot procedures
- But the accident emerges through their interaction.
- The important object is: ‘The System’

- GUDIYA makes the same argument about AI.
- We can improve:
  - models,
  - agents,
  - APIs,
  - identity,
  - cybersecurity,
  - alignment,
  - hardware safety.
- We should.
- But once they become strongly coupled at machine speed, another engineering object appears:
  - The Cognitive Field
- And that object possesses dynamics of its own.

24. A New GUDIYA Stability Principle

The F-35 investigation suggests a principle worthy of the Canon:

> *Never permit a high-consequence control-law transition solely because the system believes its operating regime has changed. Establish sufficient coherence between perceived state and externally referenced state.*

- For aviation, that problem manifested as airborne versus ground.
- For cognition it may manifest as:
  - Simulation versus Production
  - Supervised versus Autonomous
  - Observation versus Actuation
  - Normal versus Emergency
  - Isolated versus Coupled
  - Software versus Physical Authority
- The vocabulary changes.
- The systems problem does not.

25. The Deepest Lesson

- When we first watched the F-35 fall from the sky, the visible phenomenon was loss of control.
- The investigation reveals a deeper phenomenon.
- The aircraft's internal representation of its operating condition diverged from physical reality.
- Its control system then behaved according to that representation.
- That gives GUDIYA perhaps the most important lesson of the entire case:
  - A system does not have to lose intelligence to become dangerous. It can lose orientation.
  - And intelligence cannot necessarily repair orientation from inside itself.
  - Sometimes the system needs an authoritative external reference.
  - That is precisely the purpose of the Cognition Coordinate System and Grid Awareness.

Conclusion: Fly the Reality, Not Merely the Machine

The F-35 accident began far away from where it ended.

- Water contamination became a hydraulic problem.
- The hydraulic problem became a mechanical problem.
- The mechanical problem became a sensor-state problem.
- The sensor-state problem became a control-law problem.

- The control-law problem became a flight-dynamics problem.
- And corrective interventions occurred while the system itself continued changing.

That chain gives GUDIYA three enduring principles:

- Do not assume that intelligent systems know which reality they occupy.
- Do not assume that a corrective action leaves the system unchanged for the next correction.
- Do not merely predict where the system is going—predict where your intervention will send it.

- The Cognition Coordinate System tells an actant where it is.
- Cognitive Weight-on-Wheels tells it which operating regime it actually occupies.
- Cognitive State Coherence tells the Grid whether different representations of that reality agree.
- Mode Transition Confidence determines whether changing control regimes is justified.
- Intervention Trajectory Projection asks what reality a corrective action will create next.
- Together they lead to a powerful refinement of GUDIYA's Stability Engineering doctrine:

$$\boxed{Observe\ Reality \rightarrow Establish\ Orientation \rightarrow Project\ Trajectory \rightarrow}$$

$$\boxed{Project\ Intervention \rightarrow Stabilize}$$

- A pilot needs more than an intelligent airplane.
- The airplane needs more than sophisticated control laws.
- Both need trustworthy orientation to reality.
- The same will be true of machine cognition.
- **Intelligence is not orientation.
- Control is not stability.
- And the correct control law applied to the wrong reality can still destroy the system.**

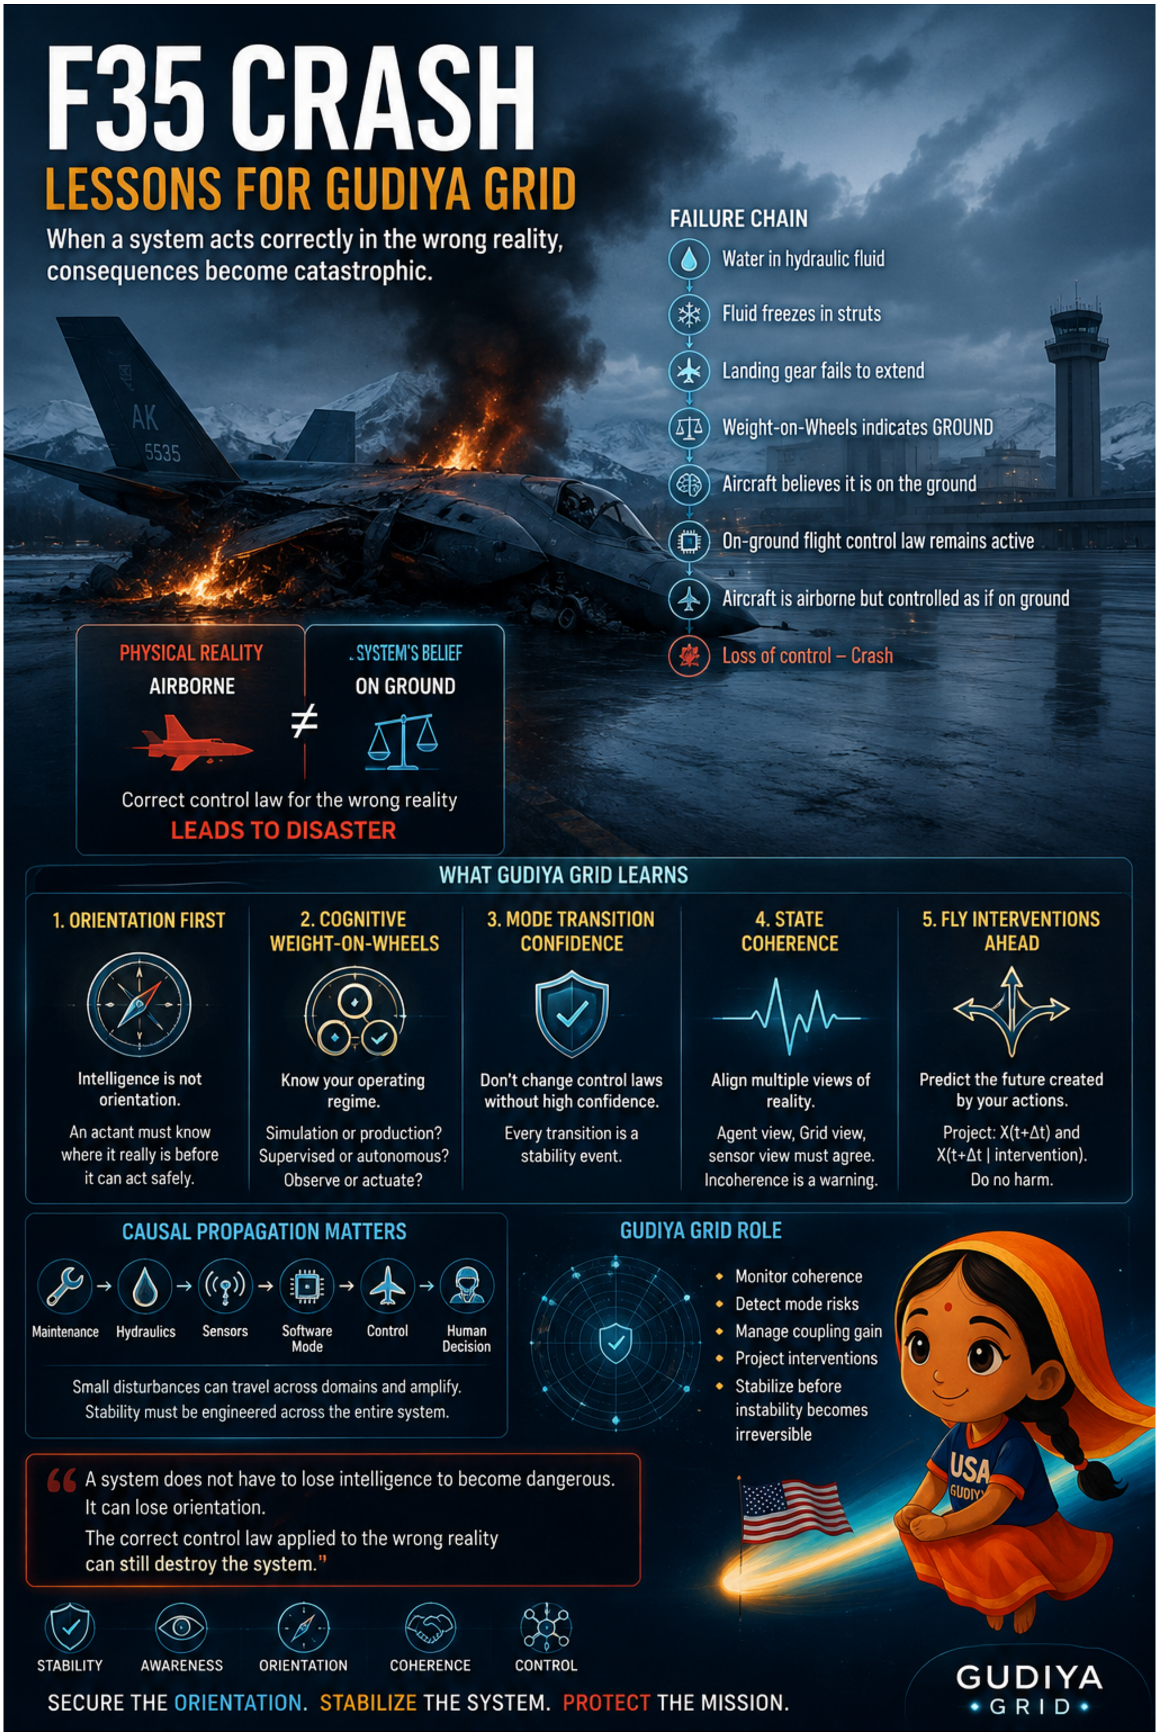
F35 CRASH
LESSONS FOR GUDIYA GRID
When a system acts correctly in the wrong reality, consequences become catastrophic.
FAILURE CHAIN
Water in hydraulic fluid
Fluid freezes in struts
Landing gear fails to extend
Weight-on-Wheels indicates GROUND
Aircraft believes it is on the ground
On-ground flight control law remains active
Aircraft is airborne but controlled as if on ground
Loss of control – Crash
PHYSICAL REALITY
AIRBORNE
≠
SYSTEM'S BELIEF
ON GROUND
Correct control law for the wrong reality
LEADS TO DISASTER
WHAT GUDIYA GRID LEARNS
1. ORIENTATION FIRST
Intelligence is not orientation.
An actant must know where it really is before it can act safely.
2. COGNITIVE WEIGHT-ON-WHEELS
Know your operating regime.
Simulation or production? Supervised or autonomous? Observe or actuate?
3. MODE TRANSITION CONFIDENCE
Don't change control laws without high confidence.
Every transition is a stability event.
4. STATE COHERENCE
Align multiple views of reality.
Agent view, Grid view, sensor view must agree. Incoherence is a warning.
5. FLY INTERVENTIONS AHEAD
Predict the future created by your actions.
Project: X(t+Δt) and X(t+Δt | intervention). Do no harm.
CAUSAL PROPAGATION MATTERS
Maintenance → Hydraulics → Sensors → Software Mode → Control → Human Decision
Small disturbances can travel across domains and amplify. Stability must be engineered across the entire system.
GUDIYA GRID ROLE
Monitor coherence
Detect mode risks
Manage coupling gain
Project interventions
Stabilize before instability becomes irreversible
"A system does not have to lose intelligence to become dangerous. It can lose orientation. The correct control law applied to the wrong reality can still destroy the system."
USA
STABILITY
AWARENESS
ORIENTATION
COHERENCE
CONTROL
SECURE THE ORIENTATION. STABILIZE THE SYSTEM. PROTECT THE MISSION.
GUDIYA GRID