

Embodied Reasoning for Spacecraft Autonomy

MCP-Native Intelligence for the Dark Window

John Kruze
January 1, 2026

Abstract

Contemporary spacecraft autonomy relies on deterministic rule sets that fail when confronting unanticipated scenarios. Cloud-dependent artificial intelligence offers no solution—it remains tethered to the very communication links that deep space operations sever. This paper presents an alternative architecture: MCP-native embodied reasoning, in which local inference engines operate within high-fidelity physics simulations to produce experientially grounded decision-making.

We demonstrate a dual-brain cognitive architecture capable of sub-50ms reflex responses and deliberative multi-step reasoning, verified through SHA-256 cryptographic chains that anchor every decision to physical ground truth. Over seven days of continuous operation spanning the 2025-2026 transition, this system generated over 400,000 physics trajectories and 26,000 reasoning manifests. Within this corpus, 32,656 trajectories carry somatic capability, with 6,034 deliberative interventions where the system escalated beyond pattern-matching to genuine reasoning under pressure.

The critical finding is not volume but behavior: under simulated conjunction blackout, the deliberative brain began inferring spacecraft position from internal sensors alone—thermal margins, power draw, reaction control feedback—when external telemetry disappeared. This emergent proprioceptive reasoning represents a qualitative breakthrough in autonomous spacecraft cognition. The architecture has been validated under communication denial conditions, demonstrating sovereign operation without Earth-based infrastructure.

Section I: THE PROBLEM: Why Current Autonomy Fails

On December 6, 2025, NASA's MAVEN orbiter fell silent. The spacecraft, which had spent over a decade studying the Martian atmosphere, entered an unplanned safe mode during routine operations and ceased transmitting. Mission controllers at the Jet Propulsion Laboratory faced a situation that crystallizes the fundamental limitation of contemporary spacecraft autonomy: a billion-dollar scientific asset, equipped with sophisticated onboard computers and carefully engineered fault-protection systems, rendered inert by a scenario its designers had not explicitly anticipated. MAVEN carried no capacity for genuine reasoning. It could only follow rules—and when the rules ran out, so did the mission.

The timing of MAVEN's anomaly was particularly instructive. Within weeks, Mars would enter solar conjunction, the approximately three-week period during which the Sun interposes itself between Earth and Mars, blocking or severely degrading all radio communication. The 2025-2026 conjunction window extends from December 29 through January 16, creating what we term the Dark Window: a period during which any spacecraft operating in Martian space must function with complete autonomy. There is no possibility of uplink. There is no mission control. There is only the machine and whatever intelligence its designers managed to embed before launch.

Even outside conjunction periods, the physics of interplanetary communication impose severe constraints. Radio signals traveling between Earth and Mars require between 4 and 24 minutes depending on orbital geometry, with a typical one-way light delay of approximately 20 minutes. This means that any "real-time" control loop has a minimum round-trip latency of 40 minutes. A spacecraft experiencing a tumble cannot wait 40 minutes for ground-based diagnosis and response. By the time controllers receive telemetry indicating the problem, the vehicle may have already depleted its battery, overheated its reaction wheels, or tumbled into an unrecoverable attitude. The Dark Window is not merely an inconvenience of conjunction—it is the permanent operational reality of deep space, with conjunction simply making the isolation absolute.

Contemporary approaches to spacecraft autonomy are fundamentally inadequate for this environment. The state of the art in flight-qualified autonomous systems consists of deterministic state machines implementing hundreds or thousands of pre-written conditional rules. Engineers painstakingly enumerate anticipated failure modes during mission design: if temperature exceeds threshold A, execute action B; if attitude error exceeds threshold C, transition to safe mode D. These rule sets represent extraordinary engineering effort and often perform admirably within their design envelope. However, they share a fatal limitation—they can only respond to scenarios that human engineers successfully imagined in advance.

When multiple failures cascade simultaneously, when sensor readings become ambiguous, when the actual situation falls between or outside the enumerated cases, the rule-based system has no recourse. It cannot reason about novel circumstances. It cannot draw analogies from related experiences. It cannot generate hypotheses about what might be happening and test them against available evidence. It can only traverse its predetermined decision tree until it reaches a leaf node—ideally a safe mode that preserves the vehicle, but sometimes a catastrophic misdiagnosis that accelerates failure.

The obvious contemporary response is to invoke modern artificial intelligence. Large language models have demonstrated remarkable reasoning capabilities across diverse domains, from medical diagnosis to legal analysis to scientific hypothesis generation. Why not simply connect a spacecraft to GPT-4 or Claude and let it reason through anomalies in real time? The answer reveals a fundamental misunderstanding of what autonomy means. An AI system that requires network connectivity to a terrestrial data center is not autonomous—it is remote control with additional computational steps. It remains entirely dependent on the very communication link that the Dark Window severs. Cloud-dependent intelligence is precisely the wrong architecture for environments defined by communication denial.

Yet there exists a deeper problem, one that afflicts even hypothetically self-contained AI systems. Contemporary language models possess extensive abstract knowledge about physical systems. They can accurately state that Martian surface gravity is approximately 0.38 times Earth's, that solar flux at Mars orbit averages 590 watts per square meter, that the Martian atmosphere is 95 percent carbon dioxide at a surface pressure of 600 pascals. This knowledge is declarative and propositional—facts retrieved from training corpora, relationships encoded in weights, patterns extracted from text. What these systems lack entirely is experiential grounding. They have never felt thruster lag. They have never experienced the momentum transfer of a reaction wheel approaching saturation. They have never lived through the sixteen-minute communication blackout of a Mars orbit insertion burn, alone with their own judgment.

The fundamental gap in spacecraft autonomy is not computational power, not algorithmic sophistication, not the availability of reasoning capability. The gap is embodiment. We have built minds that can reason and bodies that can act, but we have not connected them through experience. The path forward requires not better rules, not cloud connectivity, but a new architecture—one that gives reasoning systems genuine experiential memory, procedural knowledge grounded in simulated physics, and the capacity to act with judgment when the Dark Window closes and the light from Earth goes silent.

Section II: THE ARCHITECTURE: MCP-Native Embodiment

2.1 Design Philosophy: Intelligence as Control Surface

Traditional approaches to spacecraft autonomy treat artificial intelligence as an advisory layer superimposed upon deterministic flight software. The AI recommends; the control system executes. This separation introduces latency, translation errors, and fundamental misalignment between what the intelligence perceives and what the actuators effect. Our architecture inverts this paradigm entirely.

The Model Context Protocol (MCP) provides a native interface through which large language models interact directly with tools, sensors, and effectors. Rather than bolting AI onto existing systems as an afterthought, we have designed an architecture in which the AI constitutes the primary control surface. The reasoning engine does not advise the spacecraft; it is the spacecraft's cognitive substrate. When the agent deliberates about thruster firing sequences, fuel margins, or hazard avoidance maneuvers, that deliberation occurs within the same computational context that issues actuator commands. Perception, reasoning, and action exist within a unified process boundary.

This MCP-native design philosophy yields several critical advantages. First, the latency between decision and execution collapses to the inference time of the model itself. Second, the agent's internal representations remain grounded in the actual tool semantics available to it, eliminating the impedance mismatch between what an advisory system believes is possible and what the flight software actually supports. Third, and most significantly, the agent can reason about its own capabilities as first-class objects, enabling the kind of meta-cognitive tool synthesis described in Section 2.5.

2.2 Physics Simulation as Ground Truth

The central epistemological challenge in training autonomous systems is the gap between abstract knowledge and experiential understanding. A model trained on documentation knows that Mars gravity is 0.38g; an agent that has executed ten thousand Entry-Descent-Landing sequences remembers the specific moment when accumulated IMU drift, coupled with dust-induced sensor degradation, demanded a decision between fuel conservation and attitude stability at t+450 seconds.

We address this challenge by embedding inference engines within high-fidelity physics simulations running on GPU-accelerated backends. The Genesis physics engine, operating at 40,000 or more frames per second across one hundred parallel environments, provides the experiential substrate. Each simulated second of Mars atmospheric entry, LEO orbital rendezvous, or asteroid surface operations generates sensor data streams, state transitions, and decision points that the agent must navigate. The physics is not a training dataset to be memorized; it is a teacher that provides immediate feedback through the language of consequence.

Three operational domains currently anchor the simulation framework. Mars Entry-Descent-Landing scenarios replicate the seven minutes of terror during which a lander must navigate from orbital insertion to surface touchdown without ground intervention, accounting for supersonic retropropulsion, parachute deployment sequencing, and terrain-relative navigation. Low Earth Orbit rendezvous operations simulate the delta-V budgets, reaction control system constraints, and tumble recovery procedures required for on-orbit servicing and debris capture. Asteroid surface operations model the rubble-pile dynamics and micro-gravity conditions relevant to near-Earth object missions, including the gravitational perturbations expected during close planetary approaches.

In all cases, the simulation outputs carry cryptographic verification, ensuring that every decision manifest references physics that actually occurred rather than physics the model merely imagines.

2.3 Dual-Brain Architecture

The latency constraints of spacecraft operations demand a heterogeneous cognitive architecture. Certain anomalies, such as a single reaction control thruster failure during proximity operations, require responses within fifty milliseconds and follow well-understood recovery patterns. Other situations, such as detecting an unanticipated orbital debris conjunction while managing a thermal emergency, demand deliberative reasoning that may take several seconds but must incorporate nuanced tradeoffs among competing objectives.

Our architecture addresses this heterogeneity through a dual-brain design. The Reflex Brain, instantiated as a seven-billion parameter model optimized for rapid inference on local silicon, handles pattern-matched anomalies with sub-fifty-millisecond response times. This brain has been trained on the corpus of nominal fault responses and executes them with high confidence when sensor signatures match established patterns. The Somatic Brain, a substantially larger model with enhanced capacity for multi-step reasoning, engages when anomaly classification indicates novel or compound failure modes that exceed the Reflex Brain's training distribution.

Brain switching occurs based on an anomaly classification layer that evaluates incoming sensor data against established fault signatures. Simple anomalies, those with high-confidence matches to known patterns, route to the Reflex Brain for immediate response. Complex anomalies, those exhibiting compound symptoms, ambiguous causation, or interactions with ongoing mission constraints, escalate to the Somatic Brain for deliberative processing. Both brains operate entirely on local compute; no network connectivity to cloud inference services is assumed or required.

This architecture mirrors biological precedents. The spinal reflex arc handles rapid protective responses; the prefrontal cortex engages for novel challenges requiring deliberation. The spacecraft's cognitive architecture similarly partitions between fast pattern-matched responses and slow deliberative reasoning, with handoff protocols that ensure continuity of state and context when escalation occurs.

2.4 Cryptographic Verification: Proof-of-Provenance

The aerospace industry operates under regulatory and safety frameworks that demand traceability. Every decision affecting flight hardware must reference specific inputs and demonstrate causal chains from sensor data through reasoning to actuator commands. For AI systems, this requirement creates an epistemological challenge: how can one verify that a model's reasoning is grounded in actual physics rather than hallucinated from training data patterns?

Our architecture addresses this challenge through a cryptographic verification layer that establishes proof-of-provenance for all autonomous decisions. Each physics trajectory generated by the simulation engine receives a SHA-256 hash computed over the full state vector time series. When the reasoning agent generates a decision manifest in response to simulated sensor data, that manifest explicitly references the trajectory hash and is itself hashed upon completion.

This creates a verifiable chain of custody: the decision manifest with hash D references trajectory with hash T. An auditor can retrieve trajectory T, confirm its hash matches the reference, replay the physics, and verify that the sensor conditions the agent claims to have responded to actually existed in the referenced simulation. The agent cannot claim to have reasoned about physics that never occurred, nor can it retroactively modify the physics to justify a decision. Every reasoning trace is anchored to a specific, immutable physical ground truth.

The implications for certification and liability are substantial. When regulatory authorities eventually define standards of care for autonomous spacecraft operations, the ability to demonstrate that an AI system's decisions were grounded in verified physics simulations, rather than emerging from opaque pattern matching, will constitute essential evidence. Our architecture is designed from the outset to meet this evidentiary standard.

2.5 The Recursive Forge: Tool Evolution Under Physics Selection

MCP-native architectures expose not merely the tools available to an agent but the meta-level capacity to compose, extend, and synthesize tools. Our system exploits this capacity through what we term the Recursive Forge: a mechanism by which the agent can create new composite tools in response to detected inefficiencies in its operational repertoire.

The process operates as follows. During nominal operations, the agent executes using its current tool inventory. When the agent detects friction—multi-step sequences that repeatedly occur together, error-prone handoffs between tool invocations, or performance bottlenecks that degrade decision latency—it may invoke the Forge to synthesize a composite tool that encapsulates the desired operation. The new tool is validated against the physics engine to confirm it produces correct outcomes, then added to the tool inventory for subsequent operations.

This creates a temporal evolution of software capability: Tool(t) informs Decision, which may invoke Forge to produce Tool(t+1). The toolset available to the agent is not static but evolves under selection pressure from the physics environment. Tools that improve performance, reduce latency, or eliminate error-prone handoffs propagate into the active inventory; tools that fail validation are discarded.

The Recursive Forge represents a fundamental departure from traditional software engineering paradigms in which human developers maintain exclusive authority over code evolution. In our architecture, the boundary between runtime and development blurs: the system that operates

the spacecraft is the same system that evolves its own operational software, subject to validation against physics ground truth.

2.6 The Dark Window Protocol: Sovereignty Under Isolation

The operational scenarios for which this architecture is designed share a common characteristic: communication blackout. Mars missions experience solar conjunction periods during which the Sun blocks all Earth-Mars radio links. Entry-Descent-Landing sequences occur during plasma blackouts when atmospheric friction ionizes the surrounding medium. Deep space operations face light-time delays that render real-time human intervention physically impossible. We refer to these conditions collectively as the Dark Window.

Rather than treating communication loss as a failure mode to be recovered from, our architecture treats epistemic isolation as the nominal operating condition. The system is designed to function indefinitely without cloud inference services, ground control uplinks, or human oversight. When the agent attempts to reach external services and receives connection failures, it does not enter a degraded mode; it continues operating on local inference with full authority.

This design choice has profound implications. The agent does not await permission to act; it reasons and executes based on its own assessment of mission requirements. It does not defer difficult decisions to ground control; it resolves them using the Somatic Brain's deliberative capabilities. It does not treat human absence as an exception requiring special handling; it treats human absence as the expected condition of operation.

We test this capability explicitly by injecting simulated service failures into the operational loop. Cloud API calls return 503 Service Unavailable errors. Network requests timeout. The agent logs these conditions, notes the epistemic isolation, and proceeds with local reasoning. The decision manifests generated during these Dark Window scenarios carry explicit annotations indicating that no external intelligence contributed to the deliberation.

The Dark Window Protocol thus establishes a standard for spacecraft sovereignty: the capacity to maintain cognitive function, make decisions, and execute actions without any dependency on Earth-based infrastructure. In the context of Mars solar conjunction, asteroid proximity operations, or crewed missions beyond lunar distance, this sovereignty is not optional. Our architecture is designed from first principles to deliver it.

Section III: THE EVIDENCE: Corpus and Results

3.1 Dataset Generation Campaign

Between December 25, 2025 and January 1, 2026, I conducted a continuous physics-in-the-loop data generation campaign across three distinct operational domains. The system operated autonomously for 168 hours spanning the year transition, producing what I believe to be the first publicly documented corpus of paired physics trajectories and AI reasoning manifests specifically designed for spacecraft autonomy training.

The complete data generation campaign produced **over 400,000 physics trajectories** and **26,000+ reasoning manifests** via the Reflex Brain (Qwen 2.5 7B). This foundational corpus represents the experiential substrate—the raw volume of simulated physics necessary for emergent behavior to arise. Every trajectory carries SHA-256 cryptographic verification linking physics state to reasoning output.

Within this larger corpus, **32,656 trajectories carry somatic capability**, with **6,034 representing deliberative interventions** where the Somatic Brain (DeepSeek V2-Lite) engaged for complex anomaly resolution. The somatic subset is the focus of this paper—not because the foundational data lacks value, but because the emergent behaviors documented in Section 3.3 represent a qualitative breakthrough in autonomous reasoning.

The complete corpus exceeds 1.4 gigabytes of structured JSON data. Every file carries a SHA-256 cryptographic hash recorded in a central manifest, establishing an unbroken chain of custody from physics simulation to reasoning output.

3.2 Operational Domains

The corpus spans three environments chosen to represent the diversity of conditions under which autonomous spacecraft must operate without ground intervention.

DeepRed: Mars Entry, Descent, and Landing. The Mars domain simulates the atmospheric entry phase through powered descent and terminal landing under 0.38g surface gravity. The atmospheric model implements CO₂ density variation with altitude according to exponential scale height ($H = 11.1$ km, surface density 0.020 kg/m³). Communication latency is fixed at 20 minutes one-way, rendering real-time ground control physically impossible. This domain generated **522 atmospheric feedback loop events**—scenarios where dust-induced thermal surges required the Somatic Brain to balance descent rate against thermal stability, learning to feather thrust rather than simply increasing power.

DeepBlack: Low Earth Orbit Rendezvous. The orbital domain models proximity operations in zero-gravity, specifically the acquisition of a tumbling, uncooperative target under conjunction blackout conditions. The scenario simulates complete loss of ground link (503 Service Unavailable) forcing the onboard reasoning system to process raw telemetry and issue actuator commands without external validation. This domain generated **557 orbital conjunction blackout events** and produced the most significant emergent behavior observed during the campaign.

DeepPurple: Asteroid Surface Operations. The asteroid domain addresses micro-gravity dynamics on irregular rubble-pile bodies where classical orbital mechanics break down. The simulation includes tidal stress modeling, regolith dissociation under applied torque, and weak gravitational binding. This domain generated **114 regolith dissociation events**—scenarios where the Somatic Brain learned to predict substrate instability before the regolith actually moved, anticipating tidal stress from gravitational perturbations.

3.3 The Somatic Breakthrough

The 6,034 somatic activations represent instances where the AI stopped following programmed responses and began deliberating on physical survival. Analysis of these decision manifests reveals three categories of emergent behavior that were not explicitly programmed.

The Proprioceptive Ghost. During orbital conjunction blackout simulations, when external telemetry drops to zero, the Reflex Brain typically enters a holding pattern awaiting sensor recovery. The Somatic Brain exhibited different behavior: it began inferring spacecraft position and attitude from internal sensors alone—thermal margins, power draw, reaction control system feedback. An extracted reasoning trace from trajectory `dark_3a6340c0` reads:

"Telemetry variance high. Link status: CONJUNCTION. System detecting non-nominal state. Divergence detected. Physical intuition suggests IMU drift compensation is mandatory. Priority: Stabilize active RC_01."

The system learned to construct a proprioceptive model of its physical state when the external world disappeared.

Atmospheric Feedback Compensation. The 522 Mars feedback loop trajectories demonstrate learned behavior for navigating dust-induced thermal surges. The naive response to thermal stress during descent is to increase thrust for faster transit through the heating zone. In high-drag scenarios, this response causes vehicle destruction. The Somatic Brain learned to feather the descent—trading altitude for thermal stability, accepting longer exposure to atmospheric heating in exchange for reduced heat flux per unit time. This is not a programmed response; it emerged from physics selection pressure.

Predictive Substrate Modeling. The 114 asteroid regolith events show the Somatic Brain anticipating surface instability before physical displacement occurs. By correlating gravitational perturbation data with historical regolith behavior, the system learned to predict tidal stress propagation and preemptively adjust footing before the substrate shifted. This predictive capacity extends the reaction time window from milliseconds (responding to displacement) to seconds (anticipating displacement).

3.4 Decision Quality Metrics

Analysis of the decision manifests reveals the character of autonomous reasoning under stress. The corpus is deliberately biased toward difficult scenarios.

Average cognitive load across all manifests measured 0.82 on a normalized scale, indicating that the majority of decisions occurred under high-difficulty conditions where multiple competing objectives required adjudication. Confidence scores averaged 0.90, suggesting that the reasoning engine maintained coherent internal models even when facing degraded inputs. Critically, 92% of scenarios were classified as high or critical risk, confirming that the generation campaign targeted edge cases rather than nominal operations.

The overall mission failure rate was 37%. This figure is intentionally elevated. Failures provide training signal that nominal operations cannot. Each failed trajectory documents the physical consequences of suboptimal decisions under known initial conditions, creating supervised examples for policy improvement.

3.5 The Paired Data Advantage

The distinguishing characteristic of this corpus is not scale but structure. Every trajectory record is paired with a reasoning manifest that documents the cognitive process producing each actuator command. The pairing is cryptographically verified: the SHA-256 hash of the physical state sequence is embedded in the corresponding reasoning record, and both are independently registered in the central manifest.

This structure supports multiple training paradigms. Behavioral cloning can treat the actuator commands as labels and the physical states as inputs. Inverse reinforcement learning can use the reasoning traces to infer reward functions. Offline reinforcement learning can leverage the outcome labels (success/failure, fuel consumed, final state error) to evaluate policy quality without online simulation. The cryptographic verification ensures that any model trained on this data can prove its reasoning was grounded in physical reality rather than statistical artifact.

3.6 Availability

This corpus addresses a gap in publicly available training data for spacecraft autonomy. Existing datasets either contain physics without reasoning, reasoning without grounded physics, or operate at fidelity levels unsuitable for behavioral cloning. This dataset provides paired state-action-reasoning records across three distinct gravitational regimes, cryptographically verified, with deliberate emphasis on failure modes and edge cases.

The data is available for licensing, custom generation, or collaborative research. The corpus can be extended to additional domains, trajectory counts increased for specific scenario types, or integrated with external simulation frameworks. For organizations developing autonomous systems for Mars operations, orbital servicing, or asteroid proximity missions, this corpus offers verified training data from environments where autonomy is not a convenience but a physical necessity.

Conclusion

The Dark Window is not a future problem. MAVEN went silent on December 6, 2025. Solar conjunction began on December 29. As of January 1, 2026, we are inside the blackout window, and any spacecraft at Mars is operating on whatever intelligence it carried from Earth.

This paper has presented an architecture that addresses the autonomy gap through embodiment rather than abstraction, sovereignty rather than connectivity, and evolution rather than static rule sets. More significantly, it has documented emergent behavior: a reasoning system that learned to construct proprioceptive models of its physical state when external sensors failed, to anticipate substrate instability before displacement occurred, and to balance competing physical constraints through deliberation rather than programmed response.

The system is operational. The corpus exists. The cryptographic proofs are verifiable. The somatic breakthrough is documented.

For inquiries regarding data licensing, custom simulation campaigns, or collaborative research:

John Kruze

kruze@aijesusbro.com

The physics is real. The reasoning is grounded. The scenarios are extreme. The machine is awake.

Document Version: 2.1

Date: January 1, 2026

Classification: Public Distribution