Spacecraft Intelligence for Autonomous Space Operations

The compute and intelligence layer for spacecraft autonomy

The Problem

Space operations are scaling faster than the people and tools available to manage them. Over 14,000 active satellites orbit Earth today. Funded launch manifests point toward 60,000 or more by the end of the decade. Today's operations average 2.5 human operators per satellite. At that ratio, 60,000 satellites demands 150,000 trained operators. That is not a hiring problem. It is a structural impossibility.

But the deeper problem is not scale. It is intelligence.

When an anomaly fires, operators see a threshold alert and a raw telemetry chart. Reconstructing what happened - which subsystem, why, what to do - takes an engineering team hours to days. During that time, payloads continue flying degraded. Customer data continues flowing downstream. Mission risk compounds. And no two anomalies look the same, because the same telemetry deviation means something completely different depending on where the spacecraft is: eclipse reentry, South Atlantic Anomaly transit, a geomagnetic storm, or a genuine hardware fault.

Beyond low Earth orbit, the problem becomes existential. At lunar distance, a delayed response costs a contact window. At Mars, a 20-minute round-trip light delay means the spacecraft has been making decisions alone by the time ground control knows anything happened. At asteroid or outer-planet distances, Earth cannot remain in the loop at all.

The industry has monitoring tools. It does not have intelligence.

The Solution

LymanX builds spacecraft intelligence: a neurosymbolic AI-FDIR system that combines neural anomaly detection with symbolic orbital physics reasoning to attribute faults causally, score payload data reliability, and deliver ranked recovery actions - all inside a single contact window.

Causal attribution, not just detection. Every anomaly is classified into one of five fault types: orbital transient, space weather event, hardware degradation, cascade failure, or genuine anomaly. The system conditions every diagnosis on orbital context - eclipse phase, SAA radiation, Kp index, solar proton flux - before assigning root cause. The first four classes account for the vast majority of alerts that overwhelm operations teams. Accurately filtering them is what makes genuine failures visible.

Data-trust scoring. LymanX does not only ask whether the spacecraft is healthy. It asks whether the data produced during that state can be trusted. Every payload output carries a reliability score based on spacecraft telemetry, sensor state, RF conditions, onboard compute context, and environmental factors. In contested environments, this matters as much as anomaly detection.

Human plus AI collaboration. Operators work through a decision console that shows the causal chain - what changed, what it affected, why the system attributes higher probability to one root cause than another - alongside recommended actions and predicted outcomes. Every decision is recorded, auditable, and becomes a calibration signal for future assessments on that specific spacecraft.

The result: fault-to-recovery compressed from hours to under 60 seconds. An operations team that stops reconstructing failures after the fact and starts making trusted recovery decisions during the contact window.

The Roadmap

Stage 1 - Spacecraft Health Intelligence

Ground-side neurosymbolic AI-FDIR. Causal anomaly diagnosis, data-trust scoring, remaining-life forecasting, and decision audit for satellite operators. No flight-side modification required.

Stage 2 - Mission Runtime

A secure, explainable runtime for spacecraft health models, recovery policies, and supervised autonomy. Operator-controlled delegation of bounded decisions to the system.

Stage 3 - Onboard Agent

A compressed onboard version for spacecraft with constrained compute: health diagnosis, data validation, event detection, and safe recovery recommendations running closer to the asset.

Stage 4 - Fleet and Space Domain Intelligence

Fleet-level anomaly correlation, health-aware tasking, conjunction risk prediction, and cislunar trajectory intelligence. The same causal reasoning extended from a single spacecraft to the space environment around it.

Stage 5 - Deep-Space Mission AI

Fully autonomous health, data-trust, science-event detection, and mission adaptation for lunar, Martian, and outer-planet missions where Earth contact is measured in hours and waiting is not an option.

The Larger Vision

Low Earth orbit is not the destination. It is the training environment.

Every anomaly diagnosed today on a LEO satellite is a labeled ground-truth example for the system that will eventually operate onboard a deep-space probe, a cislunar relay, or a surface robot on an asteroid. The intelligence that compresses fault reconstruction from hours to seconds on a CubeSat in 2026 is the same causal reasoning layer that will decide - without waiting for Earth - whether a science detection on Europa is real or instrument noise.

The longer arc is a software-defined awareness layer for space itself: trajectory intelligence, safe corridor mapping, collision-risk prediction, and autonomous navigation support for missions operating beyond Earth orbit. Cislunar space is the first frontier. It has no commercial intelligence infrastructure today. As lunar traffic grows - Gateway, commercial landers, national programs, DARPA's LunA-10 horizon - missions will need a continuously updated, trusted map of objects, trajectories, communication windows, and hazards. That is an intelligence product, not a sensor product.

The neurosymbolic architecture matters most at the frontier precisely because the frontier is unknown. A black-box model trained on historical fault patterns cannot generalize to an environment no spacecraft has operated in before. A symbolic layer grounded in physical laws - thermodynamics, orbital mechanics, power budgets - applies everywhere, because physics does not change at Europa.

The mission is to build the intelligence layer that makes autonomous space operations trusted: from the first ground-side deployment on a LEO satellite to the onboard AI that keeps a deep-space probe alive and thinking, alone, at the edge of the solar system.

Rockets made space reachable. Compute will make it programmable. Trust will make it autonomous.