

Generative AI and the Future of Scientific and Human Spacecraft Operations

published by Joachim J. Kehr Editor and retired DLR/GSOC Operations Manager

The following technical assessment of future scientific and human spacecraft operations was generated by ChatGPT as prompted, reviewed, edited and concurred by JSOC Editor Joachim J. Kehr.

Executive assessment as of August 2026

Spacecraft operations are entering a transition comparable, in operational significance, to the move from operator controlled to an “embedded” generative and agentic AI control scheme. The fundamental driver is not simply the availability of better AI. It is the growing mismatch between **mission complexity and the capacity of human ground teams to supervise every decision.**

NASA's Distributed Spacecraft Autonomy program, ESA's autonomous formation-flying work, onboard AI experiments, and recent demonstrations of foundation models in orbit all point in the same direction. NASA reports, for example, that its Dynamic Targeting technology can allow an Earth-observation spacecraft to analyze imagery onboard and decide where to point its instrument within seconds. ([NASA](#)) NASA's Distributed Spacecraft Autonomy program is extending this concept from individual spacecraft to cooperating spacecraft that allocate tasks and coordinate their activities. ([NASA](#))

The new element is **generative and agentic AI**. Unlike traditional autonomy, which normally operates within a relatively narrowly defined decision space, generative AI can work with heterogeneous information—telemetry, procedures, engineering documentation, images, scientific observations, mission plans and natural-language instructions—and construct new hypotheses, explanations and candidate actions.

That makes it particularly attractive for **operations rather than merely control.**

Impact Assessment:

Area	Potential impact	State of the art	Likely operational maturity
Ground-operations assistance	Very high	Early operationalization	1–3 years
Mission planning	Very high	Demonstrations/prototypes	2–5 years
Anomaly diagnosis	Very high	Strong AI heritage + GenAI emerging	2–5 years
Scientific data triage	Very high	Already demonstrated	1–4 years
Crew assistance	Very high	Early demonstrations	2–5 years
Procedure generation	High	Prototype	2–5 years
Onboard science autonomy	Very high	Conventional AI already flying; foundation models emerging	3–8 years
Onboard mission management	High	Experimental	5–10 years
Autonomous spacecraft commanding	Very high potential / very high risk	Research	>10 years
Direct life-critical control by GenAI	Low near-term likelihood	Not mature	Probably >10 years

The key distinction is therefore:

The first spacecraft "AI co-pilots" will not be ChatGPT controlling a spacecraft. They will be tightly constrained reasoning systems embedded in a larger architecture of verified autonomy, simulation, planning and control.



Chat GPT generated image of the future DLR-GSOC ISS Columbus control room using a generative AI environment.

1. Why spacecraft operations are particularly suited to generative AI

Space operations have an unusual combination of characteristics that make generative AI attractive. The information environment is enormous.

A modern mission produces:

- telemetry;
- command histories;
- subsystem states;
- thermal and power data;
- navigation solutions;
- images;
- scientific observations;
- procedures;
- engineering documentation;
- previous anomaly reports;
- spacecraft models;
- ground-station information;

- environmental predictions.

Humans are very good at reasoning across such information, but relatively poor at doing so continuously and exhaustively. An AI system can potentially maintain a continuously updated **mission state model** and answer questions such as:

Why did the spacecraft enter this configuration?

or:

What is the most probable cause of the observed thermal excursion, and what are the three safest recovery strategies?

That is qualitatively different from conventional telemetry monitoring.

2. The application: the AI most important near-term mission control engineer

It is believed this will be the **largest near-term impact** of generative AI. Instead of replacing the flight director, spacecraft controller or subsystem specialist, an AI system becomes an extremely capable assistant. A future control-room operator might see:

Spacecraft: MAVEN

Subsystem: Power

Event: Battery temperature rising

AI assessment:

- 73% probability: degraded radiator performance
- 18% probability: sensor bias
- 6% probability: increased heater duty cycle
- 3% other

Evidence: 14 correlated telemetry signatures; similar event observed during mission days X and Y

Recommended actions:

1. Reduce non-essential payload activity.
2. Reconfigure heater group H2.
3. Request additional thermal telemetry.

Predicted consequences:

- Energy margin +4.2%
- Science loss: approximately 2.1 h
- No predicted violation of thermal constraints.

Simulation: all three candidate procedures tested against 10,000 Monte-Carlo cases.

The operator remains responsible for the decision.

This is much more plausible—and much more valuable—than giving an LLM unrestricted command authority. ESA's ORBIT-STAR work already illustrates the direction. The system demonstrated AI-based monitoring of ISS Columbus telemetry, fault detection and prediction, and identification of preventive actions. ([Europäische Weltraumorganisation](#))

3. From anomaly detection to anomaly reasoning

Traditional spacecraft autonomy generally answers: "Is something wrong?" Generative AI could increasingly answer:

"What is probably happening, why is it happening, what could happen next, and what are our options?"

This is a **major** operational step.

A conventional machine-learning system might recognise a telemetry pattern associated with a reaction-wheel anomaly.

A generative system could combine:

- current telemetry;
- historical telemetry;
- engineering documentation;
- spacecraft configuration;
- previous anomalies;
- thermal models;
- attitude history;
- recent commands;

and construct a causal explanation. The resulting architecture would be something like:

Sensors → conventional AI → state estimation → knowledge model → generative reasoning → candidate actions → verified planner → deterministic flight software

The LLM therefore becomes a **reasoning layer**, not the final actuator.

4. Scientific Spacecraft may benefit even more than communication satellites

This is one of the most interesting consequences. Traditional spacecraft operations are largely **mission-objective driven**:

“Observe target X at time Y“

Scientific autonomy can instead become **phenomenon driven**:

“Find something interesting”

NASA's Dynamic Targeting work is already demonstrating the underlying principle: a spacecraft can analyse what it sees and change its observation strategy accordingly rather than simply executing a pre-programmed sequence. ([NASA](#)). Generative AI could take this considerably further.

Imagine a solar observatory detecting an unusual magnetic structure. Instead of:

1. collect data;
2. transmit data;
3. wait for scientists;
4. scientists analyse;
5. send new commands,

the spacecraft might:

1. detect the anomaly;
2. classify it;
3. compare it with known phenomena;
4. estimate its scientific significance;

5. identify additional measurements that would discriminate between hypotheses;
6. retask instruments;
7. acquire the observations;
8. transmit the most valuable data first.

That changes the spacecraft from a **scientific instrument** into something closer to a **scientific investigator**.

5. Foundation models are already beginning to move onboard

This is no longer purely speculative.

In May 2026, NASA reported the deployment of the **Prithvi geospatial foundation model** on two in-orbit platforms, describing it as the first geospatial foundation model deployed in orbit. The model was trained on 13 years of Earth-observation data and can support multiple Earth-observation tasks. ([NASA Science](#)). This is significant because it demonstrates an important architectural transition:

AI is moving from analyzing spacecraft data after downlink → to analyzing it where the data are produced.

The consequences include:

- lower downlink requirements;
- faster response;
- autonomous event detection;
- adaptive observations;
- prioritisation of scientific data;
- reduced dependence on ground contact.

ESA has been pursuing the same broader direction for years, explicitly identifying onboard AI as a key technology for future missions. ([Europäische Weltraumorganisation](#))

6. Generative AI and autonomous navigation

Navigation is another potentially transformative area. Traditional spacecraft navigation is heavily mathematical and model based. That is unlikely to disappear. However, generative models can become useful at the **planning and interpretation layer**.

ESA-supported work on **OrbitGPT**, for example, explores foundation models capable of generating candidate spacecraft trajectories, including multi-gravity-assist trajectories and periodic orbits. ([Europäische Weltraumorganisation](#)).

The important point is not that an LLM should replace an astrodynamics solver. Rather:

Generative AI can become a search engine for the space of possible missions.

A mission designer might specify:


"Find trajectories from Earth to a near-Earth asteroid requiring less than 3.5 km/s of departure Δv , allowing one Venus flyby and providing at least six months of proximity operations."

The AI could generate promising candidates, which are then evaluated by high-fidelity astrodynamics software. This could dramatically shorten mission-design cycles.

7. The combination of GenAI + digital twins may be the real breakthrough

In my view, this is **more important than the LLM itself**.

A spacecraft AI should not reason about an abstract spacecraft. It should reason against a continuously updated **digital twin**. The architecture becomes:



Thousands or millions of simulated scenarios can be evaluated before anything is transmitted.

This is precisely the direction in which current research is moving. Digital-twin approaches are being developed for spacecraft onboard-software validation, while integrated digital/robotic twins are being investigated for rendezvous and formation-flying verification. ([MDPI](#)). The recent AI-planned Perseverance rover demonstration illustrates the principle particularly well: the generated route was checked against a digital twin before execution. ([Space](#))
The future operational chain would be:

Generate → simulate → verify → constrain → execute → observe → learn
rather than:
Generate → execute.

8. Human spacecraft operations could change even more radically

For crewed spacecraft, the value proposition is different. A crew has limited:

- time;
- attention;
- working memory;
- physical capacity;
- access to specialists;
- communication bandwidth.

The farther humans travel from Earth, the more important this becomes. On the ISS, a crew can still ask mission control. On Mars, communication delay makes real-time assistance impossible. Consequently, a Mars spacecraft or habitat will need something closer to an **embedded mission-support organization**. Generative AI could provide:

Engineering support

"Why is the oxygen-generation system consuming more power?"

Procedure assistance

"Show me the safest procedure for replacing this valve."

Scientific assistance

"What observations should we make to distinguish these two geological hypotheses?"

Planning

"Can we complete the EVA and still retain the required power margin?"

Knowledge retrieval

"Have we ever encountered this fault before?"

Training

"Simulate the likely consequences if valve V17 fails during ascent."

This could make the AI function as a combination of:

Flight controller + engineer + librarian + planner + tutor + scientific assistant.

9. The Moon and Mars are where the technology becomes strategically important

For lunar operations, communication delay is manageable but operational complexity will increase sharply. For Mars, the problem becomes fundamental. A Mars crew cannot depend on Earth for every abnormal situation. Therefore the operational hierarchy could become:

Level 0 — deterministic control

Flight computers, GNC, power protection, thermal control.

Level 1 — conventional autonomy

Fault detection, navigation, obstacle avoidance, resource management.

Level 2 — AI reasoning

Diagnosis, planning, prioritisation, scientific interpretation.

Level 3 — generative mission management

Long-horizon planning and adaptation.

Level 4 — human supervision

Crew or ground team defines objectives, constraints and acceptable risk.

This hierarchy is far more credible than the idea of a single "spacecraft AGI."

10. Swarms will amplify the impact

One spacecraft is manageable. A constellation of hundreds or thousands is fundamentally different. NASA's Starling program (not to be confused with Elon Musk's Starlink program) has demonstrated distributed spacecraft autonomy, including spacecraft communicating status and coordinating activities without continuous central control. ([NASA](#)). Generative AI could operate one level above this. Instead of telling each spacecraft:

"Execute command sequence 1247."

the operator might tell the swarm:

"Map the southern polar region with maximum spatial coverage while maintaining 20% power reserve and prioritizing transient phenomena."

The AI allocates:

- spacecraft;
- instruments;
- observation windows;
- communications;
- manoeuvres;
- data-storage resources.

The swarm becomes a **distributed instrument**. This is particularly important for future:

- interferometers;
- asteroid surveys;
- heliophysics missions;
- Earth-observation constellations;
- planetary communication networks;
- lunar navigation systems.

11. Spacecraft control centers will change

ESA's Pulse initiative is already moving toward more integrated, data-centric mission monitoring and control across missions rather than isolated mission-specific systems. ([Europäische Weltraumorganisation](#)). Generative AI could provide the missing semantic layer. Instead of operators interacting with dozens of specialist tools, the control centre might have a mission-level conversational interface:

Operator:

"Why is spacecraft B consuming more power than A?"

AI:

"The difference began 37 minutes ago following eclipse exit. The strongest correlation is heater H3. I found no evidence of battery degradation. I recommend reducing heater duty cycle from 62% to 48%. Simulation predicts no thermal constraint violation."

The underlying systems remain highly technical. The interface becomes human. That could substantially reduce operator cognitive load.

12. The most powerful application may be "mission compression"

This is a concept I consider particularly important. A conventional mission might require:

- dozens of flight controllers;
- subsystem specialists;
- science planners;
- navigation specialists;
- analysts;
- scheduling specialists;
- communications specialists.

Generative AI could allow a much smaller human team to supervise a much larger number of spacecraft. That is especially important for:

large constellations + deep space + long-duration missions.

The economic effect could be enormous. A mission that previously required continuous human supervision could become:

human-directed, AI-operated, exception-managed.

13. But generative AI has a fundamental weakness in spacecraft operations

LLMs are probabilistic generators. Spacecraft systems require:

- determinism;

- bounded behaviour;
- repeatability;
- explainability;
- fault containment;
- formal verification;
- predictable timing.

An LLM can produce a plausible but incorrect explanation. In ordinary office work this may be annoying. On a spacecraft it can be catastrophic. Therefore, **hallucination is not the only problem**. More fundamental problems include:

Distribution shift

The spacecraft encounters a state that was not represented in training.

Rare-event failure

The model performs well 99.99% of the time but fails on the 0.01% case that matters.

Non-determinism The same situation can generate different responses.

Goal misinterpretation

The AI optimises a reasonable interpretation of an ambiguous objective rather than the engineer's intended one.


Cascading errors


A wrong diagnosis produces a wrong plan, which produces a new spacecraft state, which invalidates the original diagnosis.

Adversarial or corrupted inputs spacecraft networks and command systems must assume that data can be wrong or compromised.

14. This is why "LLM directly controls spacecraft" is the wrong architecture

A credible safety architecture would look more like this:





The generative model **suggests**:

The verified autonomy system **decides whether the suggestion is permissible**.
The deterministic flight system **executes it**.

15. A fascinating intermediate technology: "explainable autonomy"

NASA's 2026 MEDOS demonstration is particularly relevant here. It uses event-driven autonomy to detect conditions and trigger appropriate spacecraft activities, while explicitly emphasising that the technology should be explainable, transparent, computationally lightweight and trustworthy. ([NASA Science](#))-

This points toward an important principle:

Generative AI should probably explain decisions made by spacecraft autonomy rather than necessarily making the final decision itself.

For example:

Autonomy:

"Entered thermal protection mode."

GenAI explanation:

"Solar illumination increased by 18% while radiator R2 was unavailable. The resulting predicted thermal margin crossed the 5°C threshold. Protection mode was therefore initiated."

This is immensely valuable for operators.

16. Scientific discovery could become autonomous

This is perhaps the most profound long-term implication. Today:

Spacecraft → data → Earth → scientists → interpretation → new observation request

Future:

Spacecraft → interpretation → hypothesis → discriminating experiment → observation → updated hypothesis

The spacecraft becomes part of the scientific reasoning loop. For example, an ocean-world probe could encounter a plume containing organic molecules. An onboard AI might recognize:

- unusual spectral signatures;
- correlations between particle populations;
- temporal changes;
- environmental context.

It could then decide: "Observation set B would provide the highest information gain for distinguishing hypotheses H1 and H2." That is effectively **autonomous experimental science**.

Research such as NASA's REASIMO/AI Space Cortex work is already exploring robust, explainable autonomy for challenging ocean-world operations where communication delays and limited mission lifetimes make human-in-the-loop responses impractical. ([arXiv](#))

17. Human–AI cooperation will be more important than AI autonomy

There is a temptation to frame the future as: Humans → AI replaces humans. For spacecraft operations, I think the more accurate model is:

Humans set objectives and constraints; AI explores the solution space; verified autonomy executes; humans retain authority over strategic decisions.

The human therefore moves **up the abstraction hierarchy**. Instead of commanding: "Turn instrument X on." the operator says:

"Maximize the probability of detecting transient atmospheric phenomena during the next 24 hours while maintaining the current energy reserve." This is a much more powerful mode of operations.

18. The human factors problem is as important as the technical one

Ironically, a very capable AI may create a new danger:

automation complacency.

If operators routinely accept AI recommendations, they may lose the ability to independently diagnose failures. Conversely, if the AI produces too many explanations and recommendations, operators may suffer from:

decision overload rather than information overload.

The ideal system therefore needs calibrated trust. It should explicitly communicate:

- confidence;
- uncertainty;
- evidence;
- assumptions;
- alternatives;
- consequences;
- what it does **not** know.

NASA's established trustworthy-AI principles explicitly emphasize safety, accountability, transparency, human-centricity and scientific/technical robustness. ([NASA Technical Reports Server](#))

19. Cybersecurity becomes much more important

Generative AI introduces a new attack surface. Consider a compromised telemetry stream. If an AI believes that: "Battery temperature is increasing rapidly" it might generate a perfectly rational response to a false premise.

Similarly, an attacker might attempt:

- prompt injection;
- corrupted operational documents;
- manipulated telemetry;
- poisoned training data;
- malicious procedures;
- model extraction;
- adversarial sensor inputs.

For spacecraft, therefore:

AI security and spacecraft cybersecurity become inseparable.

An AI should never be trusted merely because its reasoning sounds convincing.

20. Space-qualified compute is a major constraint

Large foundation models are computationally expensive. Spacecraft impose severe constraints on:

- radiation tolerance;
- power consumption;
- thermal dissipation;
- memory;
- processor availability;
- software assurance;
- communications.

This strongly favours **small, specialised models and hybrid architectures**.

Rather than putting a frontier model equivalent to a terrestrial cloud system directly onboard, spacecraft may use:

- compressed foundation models;
- quantised models;
- specialist models;
- retrieval-augmented systems;
- small language models;
- event-driven inference;
- hardware accelerators;
- conventional AI for real-time control.

Large models will probably remain primarily on Earth for some time.

21. Human missions create an additional opportunity: an onboard "institution"

This is where I see an especially profound possibility. A long-duration spacecraft or Mars habitat could contain an AI system trained not merely on general knowledge but on the **institutional knowledge of the mission**:

- spacecraft design;
- operating procedures;
- crew preferences;
- engineering history;
- previous anomalies;
- scientific objectives;
- maintenance records;
- mission rules;
- lessons learned.

Over a decade-long mission, the AI could become the spacecraft's institutional memory. This addresses a surprisingly important weakness of human organizations: People forget. Teams change. Documents become fragmented. An AI mission historian could retain the operational context continuously.

22. The "AI flight director" is plausible; the "AI astronaut commander" is not—yet

I would distinguish three levels.

Level A — AI assistant

Already technically plausible. It recommends and explains.

Level B — AI mission manager

Likely within the next decade. It can autonomously construct and execute plans within a constrained envelope.

Level C — AI commander

An AI given broad authority to decide what the mission should do under unforeseen circumstances. This requires substantially more progress in:

- robust reasoning;
- uncertainty management;
- verification;
- causal modelling;
- ethics;
- human factors;
- long-horizon planning.

I would not regard Level C as an engineering problem that has already been solved.

23. An important lesson from Proba-3

ESA's Proba-3 demonstrated autonomous formation flying with millimetre-level precision, including autonomous navigation, maneuvering, formation management and collision avoidance. ([Europäische Weltraumorganisation](#))

This is a useful reminder:

Spacecraft autonomy does not require generative AI.

Highly reliable conventional autonomy will continue to dominate domains where:

- physics is well understood;
- states are well defined;
- timing is critical;
- safety boundaries are explicit.

Generative AI adds greatest value where the problem is:

- ambiguous;
- information-rich;
- open-ended;
- difficult to encode explicitly;
- dependent on human knowledge.

That distinction should guide system architecture.

24. My projected evolution

2026–2028: AI co-pilots

Primary applications:

- telemetry summarization;
- anomaly investigation;
- procedure retrieval;
- mission planning;
- science-data triage;
- operator assistance;
- automated reporting.

Human remains firmly in the loop.

2028–2032: AI operations agents

Systems begin to:

- generate plans;
- test them against digital twins;
- coordinate multiple subsystems;
- perform routine recovery;
- optimize science schedules;
- manage communications and resources.

Human becomes **on-the-loop** for many activities.

2032–2040: Autonomous mission management

For selected missions:

- onboard foundation models;
- distributed spacecraft intelligence;
- autonomous scientific investigation;
- long-duration fault recovery;
- autonomous swarm management;
- AI-assisted crew operations.

Ground teams increasingly supervise **mission objectives rather than individual commands**.

2040+

If progress continues, the distinction between:

spacecraft, robot, scientific instrument, and mission-control system

may become increasingly blurred. A deep-space mission could effectively be an **autonomous scientific organisation distributed across spacecraft, robots and humans**.

25. Overall impact assessment

I would score the likely impact as follows:

Capability	Scientific missions	Human missions
Operations automation	★★★★★	★★★★★
Anomaly response	★★★★★	★★★★★
Mission planning	★★★★★	★★★★★
Scientific discovery	★★★★★	★★★★
Crew assistance	★★★	★★★★★
Resource management	★★★★	★★★★★
Navigation	★★★★	★★★★★
Robotics	★★★★★	★★★★★
Swarm coordination	★★★★★	★★★★
Ground staffing reduction	★★★★★	★★★★
Mission resilience	★★★★★	★★★★★
Replacement of human decision authority	★★	★★

The last line is deliberately low. **The most valuable future is not autonomous AI replacing mission controllers. It is AI allowing each human controller to supervise dramatically more complex missions.**

26. The central conclusion

The most important conclusion from the current evidence is that **generative AI is unlikely to replace conventional spacecraft autonomy; rather, it will become a higher-level cognitive layer above**

it—interpreting mission intent, synthesizing information, generating plans, diagnosing anomalies, and interacting with human operators—while deterministic and formally verified software continues to execute safety-critical control loops.

The state of the art suggests that we are approaching a **three-layer spacecraft intelligence architecture**:

1. Machine intelligence

Perception, estimation, control, navigation and fault detection. Already mature in selected applications.

2. Generative intelligence

Reasoning, explanation, planning, knowledge integration and scientific interpretation. This is the emerging frontier.

3. Human intelligence

Mission purpose, values, risk acceptance, scientific judgement and ultimate authority. This remains indispensable for high-consequence decisions. The transformative combination is therefore:

Foundation models + conventional autonomy + digital twins + verified planners + human supervision. Not: **LLM → spacecraft commands.**

The strongest evidence already exists for the first architecture: NASA is flying increasingly autonomous systems, experimenting with foundation models in orbit, developing distributed spacecraft autonomy and demonstrating AI-based science targeting; ESA is advancing autonomous spacecraft operations, formation flying, AI-based fault management and generative astrodynamics. ([NASA Science](#)).

The consequence could be profound. **Generative AI may change spacecraft operations from command execution into goal-directed supervision.** Instead of telling spacecraft *what to do*, humans will increasingly tell them *what they are trying to achieve*—with AI translating those objectives into candidate plans, conventional autonomy ensuring that those plans are safe, and digital twins providing a continuously updated rehearsal of the consequences.

For **scientific spacecraft**, this could produce the first genuinely autonomous scientific investigators.

For **human spacecraft**, it could provide something even more valuable: an onboard cognitive partner capable of carrying much of the operational knowledge of Mission Control with the crew, while allowing humans to concentrate on the decisions that actually require human judgement.

That, in my view, is the real significance of generative AI for spaceflight: **not artificial astronauts, but the emergence of spacecraft that can participate intelligently in their own operation.**