OPEN ACCESS

Electronic Informatics COMMENTARY

The Innovation Informatics

From reactive relays to proactive agents: Toward autonomous satellite networks

Ziteng Wang,[1] Yunxin Geng,[1,2,3,4] Ao Liu,[1] Zexuan Wu,[1] Ziyan Ren,[1] Baichuan Li,[1] and Weiwei Jiang[1,5,6,*]

[1]School of Information and Communication Engineering, Beijing University of Posts and Telecommunications, Beijing 100876, China

[2]State Key Laboratory of Satellite Network, Shanghai 201210, China

[3]Shanghai Key Laboratory of Satellite Network, Shanghai 201210, China

[4]Shanghai Satellite Network Research Institute Co., Ltd., Shanghai 201210, China

[5]Anhui Province Key Laboratory of Intelligent Building & Building Energy Saving, Anhui Jianzhu University, Hefei 230601, China

[6]Anhui Provincial Key Laboratory of Multimodal Cognitive Computation, Anhui University, Hefei 230601, China

*Correspondence: jww@bupt.edu.cn

Received: May 18, 2026; Accepted: November 2, 2026; Published Online: November 2, 2026; https://doi.org/10.59717/j.xinn-inform.2026.100072

© 2026 The Author(s). This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

Citation: Wang Z., Geng Y., Liu A., et al. (2026). From reactive relays to proactive agents: Toward autonomous satellite networks. *The Innovation Informatics* **2**:100072.

As 6G moves toward connected intelligence, satellites are expected to handle increasingly complex tasks onboard. Yet they operate with limited onboard resources in dynamic environments where continuous real-time ground control is difficult. Therefore, conventional intelligent strategies based on predefined decisions have limited adaptability and cannot fully exploit onboard resources. Agentic artificial intelligence (AI) offers a new direction as a task-oriented autonomous decision-making paradigm: it analyzes each task under current operating conditions and resource limits, then plans and executes accordingly, while using execution outcomes as closed-loop feedback. As shown in Figure 1, this Commentary centers on the autonomous cognition that Agentic AI brings to satellites, examines how this capability operates in satellite networks and the security issues it raises, and finally outlines the future development of agentic satellite networks.

AUTONOMOUS COGNITION DRIVING AGENTIC SATELLITE NETWORKS

Here, Agentic AI refers to a task-oriented closed-loop architecture that can autonomously analyze, decompose, execute, and reflect on a task, rather than a specific learning algorithm or multi-agent topology. When analyzing a task, Agentic AI can better identify its requirements and feed this information back into network orchestration, allowing different satellites to form a coordinated execution plan that uses their respective strengths. Security requirements can also be incorporated into this process so that autonomous decision-making and orchestration operate within defined boundaries.

Autonomous cognition

Decomposing a complex task into executable steps can allow errors to propagate across multiple reasoning stages. An evaluation and reflection stage following the Perception–Planning–Action loop can address this problem by checking intermediate execution states and triggering replanning when the external environment or the satellite state changes.[1] This process also requires memory capable of tracking changes in the constellation. Graph Retrieval-Augmented Generation (Graph-RAG) can retrieve topology-related context among satellites and ground nodes for short term memory, while Personalized Long-Term Memory (PLM) preserves historical interactions and service preferences for different user types. Together, they keep the information used for reasoning aligned with current network conditions and accumulated service experience.

Onboard model adaptation faces a different constraint. Large models may support task analysis, but repeatedly updating all model parameters is difficult under limited onboard resources. Low-Rank Adaptation (LoRA) freezes the pretrained weights and updates only low-rank parameters, greatly reducing the number of trainable parameters and the training memory burden compared with full fine-tuning. This reduces the burden of occasional onboard adaptation when a model needs to adjust to a new task or operating scenario.

Agentic reasoning does not need to handle every task. A well-defined task can be handled directly by a specialized Machine Learning (ML) model, while more complex tasks that require decomposition, contextual reasoning, or coordination invoke Agentic AI. Therefore, the satellite can decide when higher-level reasoning is worth its additional cost, balancing task quality against latency and resource use. A similar separation appears in satellite–air–ground networks, where Large Language Model (LLM)-based high-level planning works with Reinforcement Learning agents for low-level execution, reducing computational cost while retaining adaptability.[2]

Distributed orchestration

This autonomy can extend from a single satellite to satellite groups, allowing agents on different satellites to take different roles and use the spatial and functional heterogeneity across orbital layers for distributed orchestration.

The Domain Agentic System (DAS) combines reasoning agents with Service Engines for specific computational tasks.[3] When a satellite receives a task, its Agentic AI can analyze and plan the task, then invoke the appropriate local Service Engines. If the agent determines from its memory and current state that another node is better suited to handle a certain part, the agent can also reassign part of the task to another satellite. Function as a Service (FaaS) supports this process by decoupling network functions into lightweight services that can be invoked across orbital nodes, reducing deployment overhead and supporting more flexible coordination.

A hierarchical agent architecture can further map this orchestration onto the physical satellite network.[4] The Foundation model-as-agent (F-Agent) on a geostationary Earth orbit (GEO) satellite provides high level reasoning and constellation-wide guidance, while lightweight Embodied model-as-agents (E-Agents) on low Earth orbit (LEO) satellites execute specific tasks through local Service Engines. The Hybrid model-as-agent (H-Agent) connects the two levels. An E-Agent continues the current task without F-Agent intervention and requests support only when broader orchestration or coordination among satellites is needed, maintaining local execution efficiency while limiting communication overhead.

The architecture above can be illustrated by a maritime vessel monitoring task. A LEO satellite identifies potentially suspicious targets. It uses its short term and long-term memory to analyze the situation before invoking a local ML based Service Engine for rapid image processing. It also sends key processing information through the H-Agent to the F-Agent. Before the LEO leaves the area, the F-Agent uses its broader view of constellation geometry and resource states to assign another suitable LEO for continued monitoring when needed. The current LEO then passes the remaining task and a summary of its previous processing to the selected satellite through service orchestration supported by FaaS. This handover also raises trust and security questions about information sensitivity and the reliability of processing results.

Holistic trustworthiness

The autonomous cognition introduced by Agentic AI also brings new internal trust issues, and intelligence alone does not mean that an agent should be trusted for every task. Therefore, task specific trust labels can be assigned according to the task and initialized by the ground segment, then updated and verified during operation using observed behavior and satellite states. Incoming information also undergoes validity and consistency checks. During task processing, confidence thresholds can flag unreliable ML outputs, while Agentic AI reasoning can be checked and revised through feedback and updated memory. Results with low reliability can be held for further verification or handled by a validated heuristic when immediate action is required. At the constellation level, the F-Agent can use broader state information to coordinate conflicting decisions among LEO satellites and identify nodes with suspicious behavior or abnormal states. It could then adjust task orchestration or organize fault recovery. When a task is important or high risk and the

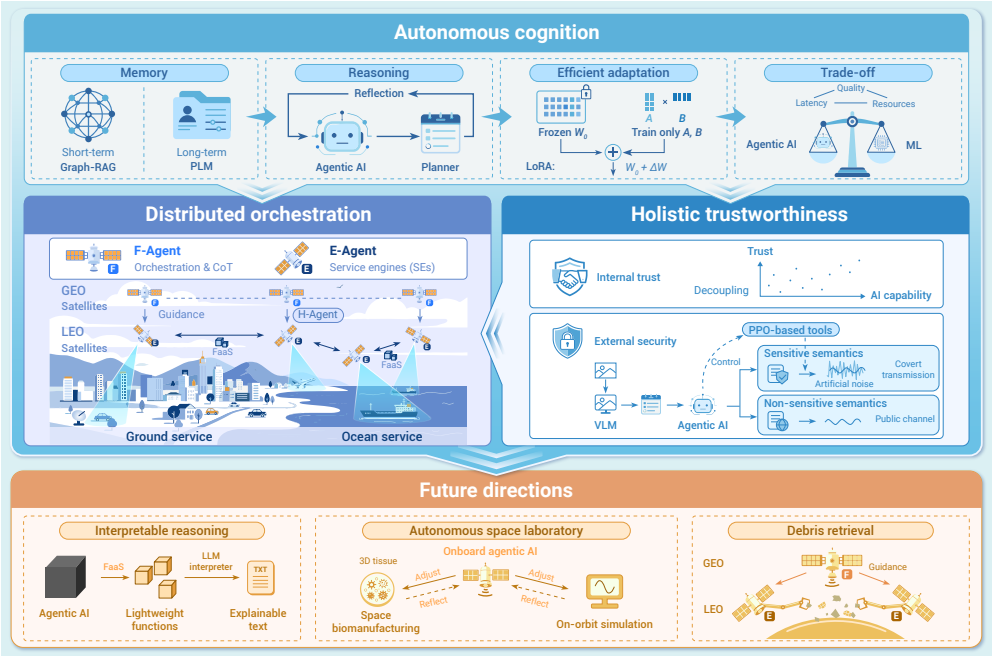


Figure 1. Framework for Agentic AI-empowered satellite networks Agentic AI places autonomous cognition at the core of satellite intelligence and extends it to hierarchical distributed orchestration. Security and trust mechanisms bound this autonomy, together supporting future satellite applications.

agent cannot make a reliable autonomous decision, confirmation from the ground is required.

Autonomous operation also places higher demands on transmission security. Agentic AI can use Vision Language Models (VLMs) to extract text-based semantics from satellite observations,[5] and then classify them using algorithms such as Proximal Policy Optimization (PPO). Non-sensitive semantics are transmitted through public channels with high transmission efficiency, while sensitive semantics are routed through covert transmission with artificial noise. This allows the system to balance transmission efficiency and security according to information sensitivity.

Taking the F-Agent described above as an example, it needs to determine whether the task requires ground confirmation and whether the selected LEO can be trusted before assignment. Then the newly assigned LEO analyzes and processes the task within the permitted action space, while continuously using reflection during task execution to check and revise its reasoning. It then protects sensitive results before transmission. For information that the LEO cannot assess with sufficient confidence, it keeps the information for further verification or requests higher-level guidance.

FUTURE DEVELOPMENT OF AGENTIC SATELLITE NETWORKS

For future Agentic AI satellite networks, an important issue is interpretability. Inspired by the FaaS architecture discussed above, the reasoning process can be divided into independent functional modules. When an error or unexpected problem occurs, the system can extract key decision parameters to analyze the cause, reflect on the reasoning process, and revise the decision accordingly. Future evaluation should examine how task oriented orchestration balances resource consumption and task completion performance, recognizing that locally suboptimal decisions may still improve end-to-end mission outcomes.

Agentic AI allows satellites to better perceive the current environment and system state, then generate actions and update them through feedback. This allows satellites to undertake tasks in space that are difficult to manage through continuous real time ground control. In space biomanufacturing, large volumes of data are generated when regulating fragile experimental environments. Onboard agents can adjust experimental conditions and operation plans according to task progress and environmental changes, reducing continuous dependence on delayed ground links. Intelligent satellites may also act as autonomous space laboratories and use actual orbital conditions, such as thermal dynamics, to conduct physical simulations in space and reduce reliance on repeated ground validation. Agentic AI can also support space debris retrieval. After detecting the position and attitude of debris, the upper layer F-Agent can provide broader coordination, while the lower layer E-Agent performs real time local analysis of the debris state before capture. This can support a more sustainable orbital environment.

Overall, Agentic AI gives satellites a degree of autonomy in space environments where continuous real time ground control is difficult. Satellites can analyze changing conditions and extend this capability to constellation level cooperation, allowing better use of orbital resources. As onboard capabilities mature, security and trust management will remain long term concerns, while stronger supervision will be needed for authority and trust issues that cannot be resolved locally.

REFERENCES

1. Jiang F., Pan C., Wang K., et al. (2026). From large AI models to agentic AI: A tutorial on future intelligent communications. *IEEE J. Sel. Areas Commun.* **44**:3507–3540. DOI:10.1109/JSAC.2026.3660010
2. Gao X., Wang Y., Liu B., et al. (2026). Agentic satellite-augmented low-altitude economy and terrestrial networks: A survey on generative approaches. *IEEE Commun. Surveys Tuts.* **28**:4800–4841. DOI:10.1109/COMST.2026.3660854
3. Antonopoulos A., Kartsakli E., Bartzoudis N., et al. (2025). AGILE-6G: Agentic AI for autonomous management of 6G network/application services. *IEEE Netw.* **2025**:1–9. DOI:10.1109/MNET.2025.3644039
4. Xiao Y., Shi G. and Zhang P. (2025). Toward agentic AI networking in 6G: A generative foundation model-as-agent approach. *IEEE Commun. Mag.* **63**:68–74. DOI:10.1109/MCOM.001.2500005
5. Feng Z., Zhang R., Qin X., et al. (2026). Toward intelligent and secure satellite communications: Agentic AI for covert semantics. *IEEE Wireless Commun.* **2026**:1–8. DOI:10.1109/MWC.2026.3666836

FUNDING AND ACKNOWLEDGMENTS

This work was supported by Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China under Grant JYB2025XDXM119; National Natural Science Foundation of China under Grants 62595741, 62595745, 62595740, 62401070, and 62321001; Natural Science Foundation of Beijing under Grant L232002 and L242012; The Open Foundation of Anhui Province Key Laboratory of Intelligent Building \& Building Energy Saving (No. IBES2025KF05ZD); The Open Project of Anhui Provincial Key Laboratory of Multimodal Cognitive Computation, Anhui University, No. MMC202403. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

Ziteng Wang conceptualized the work, reviewed the literature and existing approaches, drafted the manuscript, and prepared the figure. Yunxin Geng provided guidance on the narrative and overall manuscript and revised the language. Ao Liu contributed to language editing. Weiwei Jiang helped refine the concept and provided writing guidance and manuscript feedback. Zexuan Wu, Ziyan Ren, and Baichuan Li reviewed the manuscript and provided feedback. All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare no competing interests.