

AI-Augmented Human Decision-Making in Secure Space Operations

Sadman Saif^{1,*}, Fatemeh Sarshartehrani², Muhammad Zeeshan Karamat¹, Yahia Tawfik¹, Bo Ji², Brendan David-John² and Christiana Chamon¹

¹Bradly Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, Virginia, 24061, USA

²Department of Computer Science, Virginia Tech, Blacksburg, Virginia, 24061, USA

Abstract

This position paper envisions an integrated, AI-augmented human-machine environment designed to help operators in space (e.g., astronauts) sustain balanced workloads and stable cognitive functioning under extreme operational pressure. The proposed system combines machine-learning-based cognitive state estimation, explainable telemetry anomaly and cyber-attack detection, a large language model for guidance and assistance, and augmented-reality (AR) Heads-up Display (HUD) interfaces that mediate information flow in real time. We have developed and evaluated components, including near-real-time cognitive load monitoring and lightweight, explainable anomaly detection, to tackle the security concerns. Building on this, our long-term goal is to design critical data visualizations and AI agent support for the HUD and systematically assess their effects, advancing our broader vision of AI-supported AR interfaces in safety and security-critical environments, such as space operations.

Keywords

Augmented Reality, Cognitive State, Large Language Model, Machine Learning

1. Introduction

The expanding boundaries of space exploration expose human operators to long-duration missions that introduce a new scale of isolation. Even with extensive training, there is a growing concern for astronauts' mental health and cognitive performance [1]. Prior work has shown that as mission durations extend from 8 to 50 days, astronauts exhibit both physical and psychological impairments that degrade task performance [2, 3]. Augmented reality (AR) can enhance performance on procedural tasks in space operations (for example, carrying out maintenance on a scientific instrument) by delivering step-by-step visual guidance via a HUD [4]. However, beyond static, instruction-driven, single-task scenarios, there remains a need for explanatory data analysis, interpretation of complex system conditions, and context-aware, adaptive assistance in multitasking situations. Therefore, we envision that autonomous monitoring and maintenance systems, supported by AI and AR interfaces, can reduce astronaut workload by requiring human supervision only when necessary. Integrating AI and AR with adaptive automation alongside human operators in high-demand space environments is essential to enhance autonomy, reduce errors, and significantly improve efficiency, sustainability, response time, and overall operational effectiveness [5]. Space anomaly detection plays a crucial role in safeguarding the integrity and reliability of space systems amidst the increasing array of threats. Machine-learning methods, like ensemble ML, reinforcement learning, and deep models, can continuously monitor cyber threats such as jamming, spoofing, signal injection, and intrusions against space systems [6, 7]. In addition, Large Language Models (LLMs) function as intelligent AI assistants, particularly those designed around specific contexts like space operations, such as METIS [8] and CORE [9] facilitate smoother operations by delivering reliable outputs and multimodal feedback, thereby enhancing situational awareness and reducing cognitive workload. As reliance on AI grows in sensitive contexts such as space operations, we must

CHI'26: XR4CE Workshop, April 14, 2026, Barcelona, Spain

*Corresponding author.

✉ sadmansaif@vt.edu (S. Saif); fatemehst@vt.edu (F. Sarshartehrani); mzeeshan@vt.edu (M. Z. Karamat); mtyahia@vt.edu (Y. Tawfik); boji@vt.edu (B. Ji); bmdj@vt.edu (B. David-John); ccgarcia@vt.edu (C. Chamon)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

ensure explainability, transparent reasoning, and human oversight over all final decisions informed by autonomous components [10]. Unlike procedural AR support in [4], navigational HUD concepts mentioned in [11], and existing AI assistants such as METIS/CORE [8, 9], our specific novelty lies in integrating AI-assistance and AR into a single human-in-the-loop framework: explainable detection of telemetry security anomalies, HUD-mediated presentation of critical information, and LLM-based guidance, all coupled with real-time cognitive-load estimation. This design allows us to validate the system by measuring how AI support affects users' cognitive states in demanding multitasking scenarios.




Figure 1: AI-Augmented Human Decision-Making Envisioned Framework

In space operations, we argue that human-AI teaming should prioritize cognitive well-being and calibrated autonomy over maximal automation. This position paper presents our vision and initial prototyping efforts of an AI-augmented environment that integrates cognitive state estimation, explainable anomaly monitoring, a domain-specialized LLM assistant, and AR heads-up interfaces to support astronauts under extreme demands. We detail our framework, report progress on a working prototype, and outline a research agenda for evaluating AI-mediated cognitive support in safety-critical settings. Our work advances transparent, human-centered AI systems and invites interdisciplinary feedback to refine interface design, evaluation methods, and ethical safeguards.

2. AI-Augmented Environment for Space Operation

Our current prototype framework is integrated within a Unity-based VR environment on the International Space Station to simulate a space-themed AR environment with AI-augmented components. Multiple components must function concurrently to monitor, alert, and assist astronauts, focusing their attention on the most critical information. This helps them make the best decisions with minimal stress and full attention. Figure 1 shows the envisioned framework in its entirety. This envisioned system integrates:

- Real-time cognitive state monitoring
- Cyber-attack detection from telemetry data
- LLM-based AI-assistant
- HUD for real-time data visualization

We evaluate a cognitively demanding, safety-critical spacecraft control task in which participants start off by setting spacecraft velocity based on distance and target arrival time. While doing this initial task,

a simulated sequence of telemetry alerts will be presented through the HUD. To induce high cognitive load, the simulation includes true anomalies, false negatives, and near-threshold normal cases that may be incorrectly flagged by the telemetry attack detection model; participants must inspect 3–4 feature values and judge whether each alert reflects a true anomaly or a false flag, without knowing the ground truth. They will complete half of these cases using only the pre-simulation briefing and the other half with explainable AI guidance delivered through the AR HUD, while cognitive load is monitored across conditions. In the final phase, participants will correct true telemetry anomalies by calculating and updating the appropriate feature values to restore spacecraft navigation, with voice-based assistance from an LLM agent for procedural guidance and mathematical calculation support. This will allow us to assess whether dynamic AI assistance reduces workload and improves performance in challenging environments. We will assess effectiveness using task performance and subjective workload (NASA-TLX), examining changes in cognitive load across conditions.

2.1. Real-time cognitive state monitoring system

The effectiveness of the system is validated by observing how the astronaut’s cognitive state responds to changing situations, the support from the AI agent, the interpretability of ML models’ outputs, and the use of augmented reality environments. In [4], cognitive workload is assessed after task completion only using NASA-TLX questionnaire, where participants self-report their perceived workload. This retrospective approach may introduce bias or inaccurate recall of the actual cognitive demand experienced during the task. In contrast, real-time cognitive load monitoring with trained ML models can provide a more objective and continuous measure, enabling stronger validation of the impact of AI-assisted systems. For real-time monitoring of cognitive load, eye-tracking data is well-suited for this context, as it is headset-integrated, whereas other physiological data like Electroencephalography, Electrocardiography, or Electrodermal Activity, require additional electrodes and complex data processing [12]. In recent works on cognitive load observation, Tobii Fusion Pro and HTC Vive Pro headsets are used, which can extract real-time eye-tracking data with a good user experience [13, 14]. Our framework uses the Meta Quest Pro to extract users’ eye-tracking data in near real time, with a 14 ms delay, to assess their cognitive state and update it on the HUD. Continuous cognitive-state assessment can verify how other components reduce astronauts’ stress and workload. Although designing ecologically valid experiments that realistically mimic cognitive and physical stressors of microgravity is challenging, it is essential to demonstrate the substantial impact of AI assistance. With the dataset mentioned in [14], we have established near-real-time eye-tracking data processing and prediction using an LSTM model, with a 96% accuracy across standard metrics.

2.2. Cyber-attack Detection System

Among different cyber attacks, we have developed a lightweight telemetry attack detection model using the OPS-SAT dataset [15], which has several orders of magnitude smaller model size than other Fully Connected Neural Network and transformer-based models [16]. The use of this lightweight model results in only a 1.1%–3.4% reduction in performance metrics, making it well-suited for resource-limited environments such as space operations. The model uses explainable AI to highlight key features behind anomaly alerts and provide more detail when needed. As part of an autonomous monitoring and detection system, it sends alerts to astronauts via their HUD, along with an analysis explaining the features that drove the model’s decision. This aligns with recent HCI work on calibrated trust in AI systems, where explanations support human verification and intervention.

2.3. LLM as an AI-assistant

The proposed framework integrates a domain-specialized Large Language Model (LLM) to support astronauts during routine and procedural operations. The model is explicitly constrained and grounded through fine-tuning on verified mission documentation, including ISS operational procedures, maintenance manuals, and safety checklists. Astronauts interact via voice input and receive concise, context-

aware guidance through a HUD, enabling rapid access to certified information without manually consulting extensive documentation, thereby reducing cognitive load while preserving the astronaut's role as primary decision-maker. A key challenge is ensuring reliable, verifiable LLM guidance while preventing hallucinations and resisting adversarial, prompt-based, and distributional threats. To mitigate automation bias, the AI operates strictly as a decision-support system. It continuously communicates predictions and explanations via the HUD and defers to a human-in-the-loop framework under ambiguous or anomalous conditions.

2.4. Heads-up Display (HUD) and Augmented reality (AR) environments

Design of a pleasant and helpful HUD interface is as important as providing the critical information to assist the users. In [11], the authors provide core design insights for astronaut HUDs, emphasizing the need to reduce visual and cognitive overload while balancing strong navigational guidance with preservation of astronaut autonomy. Its findings show that although directive HUD support can improve usability and lower workload, overly prescriptive guidance may encourage overreliance and constrain independent situational assessment. Keeping those insights in mind for interface designing, our position paper adds that in high-pressure, multitasking scenarios, AI can further enhance human-machine interaction by providing adaptive, context-aware support that reduces cognitive burden, while keeping final analysis and decision-making under human control. All integrated components communicate user information through the HUD, including cognitive status, anomaly alerts, explanations, critical details, and AI agent assistance. This interface prioritizes safety-critical anomalies and cognitive overload indicators over routine work. Beyond HUD, VR/AR technologies have the potential to facilitate relaxation and mitigate feelings of isolation by simulating earth-like environments during times of need [17]. When cognitive load and system state permit, the system can suggest opportune times for users to enter these relaxation environments, while still surfacing critical alarms via the HUD.

2.5. Current prototype

The current prototype framework, which emphasizes integrating the various components, is illustrated in Fig. 2. This is built in Unity for the Meta Quest Pro that hosts a VR space station setting to emulate augmented reality conditions for space operations. Cognitive load is tracked in near real time via eye-tracking data, and the headset's HUD presents the output of the cognitive load detection model. In parallel, a telemetry monitoring system performs anomaly alert processing. The user initially engages with the environment to complete the first task depicted in Fig. 2(a). In the case of an alert, the user can then choose to view the specific telemetry features that triggered the anomaly, as shown in Fig. 2(b). After completing the component integration, our next step is to refine the AR interactions via the HUD to make them more polished, intuitive, and user-friendly based on the feedback of pilot studies.

3. Future Research Directions

In the near term, we plan (1) a formal user study to evaluate the current prototype's impact on workload and decision quality, and (2) co-design sessions with space operations experts to ground and extend our design space. Beyond that, emerging open research questions include the following:

- **Feedback mechanism:** How can we design a feedback mechanism that enables astronauts to effectively retrain when AI assistance is misaligned with the true system state?
- **Extension of the use of AI:** To what extent can monitoring and routine operational tasks be further automated using AI, and where should clear boundaries be drawn?
- **Resource constraints:** Considering computational, memory, and energy resource constraints in spaceborne platforms, what levels of model complexity are feasible without compromising reliability, safety, and real-time performance?




Figure 2: (a) In the first scenario, the user is interacting with the velocity setting interface in the environment while getting update of the anomaly detection and cognitive load. (b) In the case of an alert, the user can see the most important features about the anomaly and respond based on the user's analysis.

4. Conclusion

We believe that comprehensive integration of this framework will create a human-AI collaborative synergy that facilitates a well-informed decision-making process. Live, interactive AI assistance with AR interfaces can significantly contribute to maintaining user cognitive well-being, particularly in high-stakes operational contexts such as space missions. Our framework provides an initial step towards this vision.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT-4 and Grammarly for grammar and spelling checks, as well as for paraphrasing and rewording. After using these tool(s)/service(s), the author(s) reviewed and edited as needed and take(s) full responsibility for the publication's content.

References

- [1] F. Pagnini, Supporting the mind in space: Psychological tools for long-duration missions, *Interact J Med Res* 13 (2024) e66626. URL: <https://www.i-jmr.org/2024/1/e66626>. doi:10.2196/66626.
- [2] A. Arone, T. Ivaldi, K. Loganovsky, S. Palermo, E. Parra, W. Flamini, D. Marazziti, The burden of space exploration on the mental health of astronauts: A narrative review, *Clinical Neuropsychiatry* 18 (2021) 237–246. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8696290>. doi:10.36131/cnfioritieditore20210502.
- [3] E. Takács, I. Barkaszi, I. Czigler, L. G. Pató, A. Altbäcker, J. McIntyre, G. Cheron, L. Balázs, Persistent deterioration of visuospatial performance in spaceflight, *Scientific Reports* 11 (2021) 9590. URL: <https://www.nature.com/articles/s41598-021-88938-6>. doi:10.1038/s41598-021-88938-6.
- [4] A. M. Braly, B. Nuernberger, S. Y. Kim, Augmented reality improves procedural work on an

- international space station science instrument, *Human Factors* 61 (2019) 866–878. doi:10.1177/0018720818824464.
- [5] K. Thangavel, R. Sabatini, A. Gardi, K. Ranasinghe, S. Hilton, P. Servidia, D. Spiller, Artificial intelligence for trusted autonomous satellite operations, *Progress in Aerospace Sciences* 144 (2024) 100960. URL: <https://www.sciencedirect.com/science/article/pii/S0376042123000763>. doi:<https://doi.org/10.1016/j.paerosci.2023.100960>.
 - [6] A. Fejjari, A. Delavault, R. Camilleri, G. Valentino, A review of anomaly detection in spacecraft telemetry data, *Applied Sciences* 15 (2025). URL: <https://www.mdpi.com/2076-3417/15/10/5653>. doi:10.3390/app15105653.
 - [7] A. Diro, S. Kaisar, A. V. Vasilakos, A. Anwar, A. Nasirian, G. Olani, Anomaly detection for space information networks: A survey of challenges, techniques, and future directions, *Computers Security* 139 (2024) 103705. URL: <https://www.sciencedirect.com/science/article/pii/S0167404824000063>. doi:<https://doi.org/10.1016/j.cose.2024.103705>.
 - [8] C. Hartmann, F. Speth, D. Sabath, F. Sellmaier, Metis: An ai assistant enabling autonomous spacecraft operations for human exploration missions, in: *2024 IEEE Aerospace Conference*, 2024, pp. 1–22. doi:10.1109/AERO58975.2024.10521154.
 - [9] L. Bensch, O. Bensch, T. Nilsson, Advancing Intelligent Personal Assistants for Human Spaceflight, in: L. Bensch, T. Nilsson, M. Nisser, P. Pataranutaporn, A. Schmidt, V. Sumini (Eds.), *Advancing Human-Computer Interaction for Space Exploration (SpaceCHI 2025)*, volume 130 of *Open Access Series in Informatics (OASIs)*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2025, pp. 18:1–18:18. URL: <https://drops.dagstuhl.de/entities/document/10.4230/OASIs.SpaceCHI.2025.18>. doi:10.4230/OASIs.SpaceCHI.2025.18.
 - [10] L. Bensch, O. Bensch, T. Nilsson, J. A. Paradiso, S. Davidoff, L. Blanchard, P. Pataranutaporn, M. Weber, J. N. Kutz, F. Flemisch, Human-in/on-the-Loop AI: Enabling Human Controllability and Decision-Making in Spaceflight, *Springer Nature Singapore*, Singapore, 2025, pp. 1–46. URL: https://doi.org/10.1007/978-981-97-8440-0_40-1. doi:10.1007/978-981-97-8440-0_40-1.
 - [11] L. Bensch, T. Nilsson, J. Wulkop, P. Demedeiros, N. D. Herzberger, M. Preutenborbeck, A. Gerndt, F. Flemisch, F. Dufresne, G. Albuquerque, A. Cowley, Designing for human operations on the moon: Challenges and opportunities of navigational hud interfaces, in: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, Association for Computing Machinery, New York, NY, USA, 2024. URL: <https://doi.org/10.1145/3613904.3642859>. doi:10.1145/3613904.3642859.
 - [12] Q. Li, O. Molloy, H. El-Fiqi, G. Eves, Applications of machine learning in assessing cognitive load of uncrewed aerial system operators and in enhancing training: A systematic review, *Drones* 9 (2025). URL: <https://www.mdpi.com/2504-446X/9/11/760>. doi:10.3390/drones9110760.
 - [13] E. Ktistakis, V. Skaramagkas, D. Manousos, N. S. Tachos, E. Tripoliti, D. I. Fotiadis, M. Tsiknakis, Colet: A dataset for cognitive workload estimation based on eye-tracking, *Computer Methods and Programs in Biomedicine* 224 (2022) 106989. URL: <https://www.sciencedirect.com/science/article/pii/S0169260722003716>. doi:<https://doi.org/10.1016/j.cmpb.2022.106989>.
 - [14] J. N. Setu, J. M. Le, R. K. Kundu, B. Giesbrecht, T. Höllerer, K. A. Hoque, K. Desai, J. Quarles, Mazed and confused: A dataset of cybersickness, working memory, mental load, physical load, and attention during a real walking task in vr, 2024. URL: <https://arxiv.org/abs/2409.06898>. arXiv:2409.06898.
 - [15] B. Ruszczak, K. Kotowski, D. Evans, J. Nalepa, The ops-sat benchmark for detecting anomalies in satellite telemetry, *Scientific Data* 12 (2024). doi:10.1038/s41597-025-05035-3.
 - [16] A. Fejjari, A. Delavault, R. Camilleri, G. Valentino, Transformer-based anomaly detection for satellite telemetry data, *Acta Astronautica* 238, Part A (2026) 739–745. doi:10.1016/j.actaastro.2025.09.0.
 - [17] S. Holt, Virtual reality, augmented reality and mixed reality: For astronaut mental health; and space tourism, education and outreach, *Acta Astronautica* 203 (2023) 436–446. URL: <https://www.sciencedirect.com/science/article/pii/S0094576522006889>. doi:<https://doi.org/10.1016/j.actaastro.2022.12.016>.