Current AI Technology in Space

Justin Goodwill, Christopher Wilson, James MacKinnon
NASA Goddard Space Flight Center
8800 Greenbelt Rd, Greenbelt, MD, 20771
{justin.goodwill, christopher.m.wilson, james.mackinnon}@nasa.gov

1. Introduction

Across the science and defense space sectors, researchers and spacecraft designers are actively exploring the use of artificial intelligence (AI) and machine learning (ML) onboard spacecraft for applications requiring low-latency predictions and/or data products to enable next-generation mission concepts. However, there are many modeling and computational challenges in attempting to deploy AI onboard spacecraft.

From a modeling standpoint, the space domain must address many of the same challenges as the commercial embedded domain for mobile and Internet of Things (IoT) deployment. Commonly used supervised AI models, especially those with deep architectures, depend upon vast amounts of labeled training data to extract rich, meaningful features and accurately predict the output. However, for novel instrument sensors or missions to unexplored environments, there is often a lack of large-scale datasets that accurately represent the full distribution of input data. Thus, it can be challenging to ensure the model generalizes well to new data captured by the instrument sensor in its planned environment.

From a computational standpoint, there are numerous challenges for spacecraft processors due to constraints imposed by the harsh space environment. The radiation environment in space can cause issues ranging from non-destructive soft errors to even catastrophic device failures from varied mechanisms, including cumulative total ionizing dose (TID) effects and single event effects (SEEs). Beyond radiation, spacecraft designers are limited by further constraints, including the size, weight, power, and cost (SWaP-C) requirements of their target mission. While current radiation-hardened (rad-hard) processors provide the necessary resiliency for space radiation effects, they are extremely performance-limited, trailing generations behind commercial embedded processors. These performance limitations render the deployment of state-of-the-art AI frameworks on rad-hard platforms generally infeasible. As such, an attractive option is the use of commercial-off-the-shelf (COTS) embedded devices, ranging from CPUs, GPUs, field-programmable gate arrays (FPGAs), and custom accelerator application-specific integrated circuits (ASICs). However, it is essential that these commercial parts are screened for space use, which includes radiation testing to examine their response to TID and SEEs. Once screened, various fault-tolerant computing techniques can be employed to enhance their reliability in the presence of radiation-induced bit upsets. Given their massive performance increase over their rad-hard alternatives, these COTS devices offer promising solutions to feasibly deploying AI models onboard spacecraft.

2. Onboard AI Applications

The use of AI onboard satellites is driven by the need for low-latency predictions or products for a variety of autonomous applications. Historically, there have been many demonstrations of autonomous functionality, but due to limitations in onboard processing capability, the machine-learning and computer-vision algorithms implemented have been fairly simple with respect to computational complexity.

Most early demonstrations of autonomy involved technology-demonstration missions to increase the technology readiness level (TRL), a measure of its technological maturity, such that these concepts could eventually be used in higher class and potentially flagship missions. As part of the New Millennium program to advance technology for land-imaging instruments, Earth Observing-1 (EO-1) began its Autonomous Sciencecraft Experiment in 2003, which deployed several onboard autonomous remote-sensing applications, including cloud detection, flood scene classification, change detection, and generalized feature detection, using hyperspectral data from the Hyperion instrument. However, the complexity of the algorithms deployed was severely limited by the computational capability of the onboard Mongoose V microprocessor, primarily featuring simple decision tree and thresholding techniques [1]. Approximately a decade later in 2013, a 1U ($10 \times 10 \times 10$ cm^3) technology-demonstration CubeSat, known as the Intelligent Payload Experiment (IPEX), advanced these autonomous onboard processing concepts, using more complex random forest algorithms for image classification and salience-map algorithms to identify "interesting" regions in imagery for downlink [2].

More recently, higher class missions have integrated autonomous optical navigation due to the infeasibility of using human-in-the-loop control for certain spacecraft maneuvers. To perform the Touch-And-Go maneuver on the asteroid Bennu's surface, the Origins, Spectral Interpretation, Resource Identification, Security, Regolith Explorer (OSIRIS-Rex) spacecraft employed a natural feature-tracking algorithm. Using a database of features collected during a flyby of Bennu, the onboard processor performed cross-correlation algorithms to autonomously identify these features in images captured during the Touch-and-Go maneuver and estimate the spacecraft's position and contact point [3]. Similarly, to improve landing accuracy and avoid hazards (e.g., rocks, craters, steep slopes, sand ripple fields, etc.) in Jezero Crater, the Mars 2020 science mission employed autonomous terrain relative navigation (TRN) algorithms during the entry, descent, and landing (EDL) phase of the mission. These algorithms correlated features from camera pictures taken during the descent with an onboard hazard map, which is generated *a priori* using data products from the Mars Reconnaissance Orbiter, to simultaneously determine the spacecraft's location/altitude and navigate to safe zones [4].

With advances in onboard computing (discussed further in the Section 3), spacecraft designers are now investigating the deployment of more computationally complex AI models due to their performance over the last decade on computer-vision and natural-language-processing tasks. Beyond the aforementioned historical examples, Table 1 lists key domains and applications that could benefit immensely from the use of onboard AI for low-latency predictions.

Domain	Applications
Remote Sensing	• Rapid Disaster Response (e.g., Wildfire Detection) • Data Triage including Image and Video Compression • Onboard Product Generation
Guidance, Navigation, and Control (GNC)	• Autonomous Rover Controls • Autonomous Hazard Detection and Landing • Horizon/Star Tracking • Terrain Classification
Mission Planning	• Intelligent Scheduling • Distributed System Missions
Communication	• Software Defined Radio • Cryptography

Table 1: Domains and applications that could significantly benefit from the use of onboard AI

While an exhaustive review of onboard AI examples is beyond the scope of this chapter, we will highlight major domains where onboard AI research is bourgeoning, including remote sensing and autonomous guidance, navigation, and control (GNC). For remote-sensing applications, methods for onboard data triage are becoming exceptionally critical to maximize the science return of remote-sensing instruments, as limited satellite downlink bandwidth becomes saturated by ever-increasing amounts of data generated by next-generation sensors. Currently, numerous missions spanning planetary science, Earth science, and other domains suffer from severe downlink restrictions. For instance, the four Magnetosphere Multiscale (MMS) spacecraft generates approximately 100 GB/day of science data in its high data rate mode, but only 4% of this data can be transmitted to the ground on average [5]. AI models can be used to classify objects in an image or identify regions of interest, allowing the spacecraft to autonomously prioritize what data to downlink and/or track interesting features for capture. For example, for Earth observation, onboard cloud-masking algorithms could segment an image with cloudy and non-cloudy labels. Since clouds constitute a substantial portion of the imagery captured, but are irrelevant to many missions, onboard cloud masking can significantly reduce the data volume that must be downlinked. In particular, the Φ-sat-1 experiment was one of the first technology-demonstration satellites to incorporate and test this onboard cloud-masking capability, using a deep convolutional neural network (CNN), known as CloudScout, for cloud detection on imagery captured by a hyperspectral camera, HyperScout-2. With extremely limited preflight data captured by the HyperScout-2, the team initially trained the model using a proxy dataset generated by augmenting existing Sentinel-2 datasets [6]. Numerous other AI models have been proposed for cloud masking on different sensors, which can be deployed onboard assuming sufficient computational resources are available.

Similarly, the ability to provide low-latency classifications of Earth imagery is also critical for rapid disaster response in the remote-sensing domain. Recently, the deployment of large constellations of small satellites, such as Planet Lab's Dove constellation [7], offers unprecedented spatial and temporal coverage of Earth compared to flagship, single-spacecraft missions, significantly reducing revisit times in disaster areas. For natural disaster detection via these remote-sensing satellites, the detection time is limited by lengthy downlink times of the raw data due to limited bandwidth and by ground processing time. To reduce these detection times, researchers are transitioning towards more onboard processing such that natural disasters can be detected in near-real time, with alerts/warnings able to be downlinked rapidly as a significantly smaller fraction of the broader raw sensor data. Among natural disasters, wildfire detection and flood mapping are two applications that have garnered significant research interest, with many groups examining various AI models for accurate detection and segmentation of the affected areas [8]-[12]. A key consideration for onboard deployment is designing the architecture of the neural-network model such that it can be feasibly deployed on the spacecraft's computational hardware. Notably, the authors of [11] specifically crafted their model to be efficiently executed on the computational hardware present on Φ-sat-1, demonstrating the segmentation of a 12-megapixel image in less than a minute.

With respect to the autonomous GNC domain, major strides have been made toward building large-scale datasets necessary to train AI models for planetary rovers and landers. While multiple Mars Rovers, including Spirit, Opportunity, Curiosity, and Perseverance, have employed an autonomous-driving capability, known as AutoNav, even the most advanced versions of their machine-vision processing pipeline have been constructed purely on classical computer-vision algorithms, relying solely on geometric information to traverse the Mars landscape [13]. However, with Spirit and Curiosity having been stuck in sandy terrain and Curiosity's wheels being punctured on sharp rocks, there is a need to autonomously identify semantic information about the terrain type to assess traversability of the landscape, similar to how self-driving cars employ semantic-segmentation models to identify the drivable surface. As such, for future Mars rover missions, NASA's Jet Propulsion Lab (JPL) has developed the AI4Mars dataset that includes approximately 326K semantic-segmentation instances (among four classes: "soil," "bedrock," "sand," and "big rock") in 35K images from the Spirit, Opportunity, and Curiosity rovers in order to train common deep-learning models for semantic segmentation [14]. For autonomous-landing capabilities, the Mars 2020 TRN algorithms relied primarily on classical computer-vision techniques based on template matching and registration to *a prior*i hazard maps. For relatively unmapped and dynamic environments such as Europa, these TRN techniques may be infeasible, as they are heavily dependent on *a priori* hazard maps. Researchers are thus attempting to adapt AI models from the autonomous-driving domain to achieve much more generalized perception, using models such as You Only Crash Once (YOCO), which was first trained on simulated data and then tested on real Mars 2020 data, to predict both the location and semantic information of hazards during landing [15].

However, unlike Mars 2020, many science missions employ highly specialized sensors that capture novel, first-time measurements, which poses several modeling challenges for AI applications. Notably, the process of capturing sufficient quantities of this novel data for model training can be prohibitively restrictive, and labeling the data is often time-consuming that may require input from subject-matter experts. The combination of these factors often leads to the lack of a large-scale labeled training dataset for a science instrument that may not capture the full distribution of input data seen during deployment, as the Φ-sat-1 developers experienced when developing the cloud-detection model for HyperScout-2. Certainly, proxy data from similar sensors or physics-based simulations can be used to initially train the model, as the Φ-sat-1 developers have done. Still, with limited data in the planned deployment environment, it can be challenging to validate the generalization of the AI model. As such, during deployment, there is a need to implement some form of continuous validation, especially for mission-critical and safety-critical applications, in order to trust the model's predictions. Continuous validation involves the downlinking of raw data, with the quantity depending upon the performance requirements of the model. If significant data-mismatch error or model drift is encountered, re-training and then updating the model onboard can also be challenging depending upon the model size, as uplink rates are often even more limited than downlink rates.

3. Space Computing Devices

Deploying AI models onboard spacecraft is challenging because common rad-hard space processors are extremely performance-limited, lagging generations behind commercial embedded technology. Because they are expensive to develop and serve vital science/defense goals, flagship missions have traditionally used slow, rad-hard processors to ensure reliability and safeguard against failure due to radiation effects. Table 2 lists the most prevalent radiation-hardened processors and examples of the high-class missions for which they have been used.

Radiation-Hardened Processor	Description	Missions Served
BAE RAD6000 [16]	Radiation-hardened 32-bit CPU, employing the Power Architecture of the IBM RISC System/6000	Mars Exploration Rovers (Spirit/Opportunity), Deep Space 1, Spitzer Telescope, DSCOVR
BAE RAD750 [17]	Successor to the RAD6000, employing the radiation-hardened version of the IBM PowerPC 750	Deep Impact, Mars Reconnaissance Orbiter, Curiosity and Perseverance Rovers, James Webb Space Telescope
CAES Gaisler GR712RC [18]	Dual-core processor implementing the LEON3FT, a fault-tolerant version of the LEON3 SPARC V8 processor	DART/LICIAcube, Artemis I/ArgoMoon
CAES Gaisler GR740 [19]	Successor to the GR712, featuring a quad-core LEON4FT, a fault-tolerant version of the LEON4 SPARC V8 processor	Copernicus, Nancy Grace Roman Space Telescope (Planned)
BAE RAD5545 [20]	radiation-hardened, quad-core processor implementing Freescale Semiconductor's PowerPC e5500v architecture	Lunar Gateway Power and Propulsion Element (PPE) (Planned)

Table 2: List of common rad-hard processors

The process of radiation hardening such complex CPU devices necessitates considerable financial and engineering investment, with many rounds of radiation testing at particle-beam facilities. Once in production, strict quality assurance measures must be implemented during the fabrication process to achieve Qualified Military Line (QML) classification. The combination of radiation-hardening design and extensive development cycles often leads to radiation-hardened CPUs using archaic architectures and older fabrication processes with larger feature sizes, making them slower (clocked at lower frequencies) and more power hungry than commercial embedded technology. To demonstrate the gap between rad-hard and commercial technology, Figure 1 compares the computational density (CD), a metric that measures the theoretical steady-state performance of the computational units of a processor for a stream of independent integer (Int8, Int16, Int32) and floating-point (SPFP-Single Precision Floating Point, DPFP-Double Precision Floating Point) operations, of state-of-the-art rad-hard processors (BAE RAD750, CAES Gaiser GR740, and BAE RAD5545) and commercial embedded processors (Xilinx Zynq 7020 and Intel Core i7-4610Y) in millions of operations per second (MOPS) on a logarithmic scale [21]. Even the latest rad-hard processor, the BAE RAD5545 (still awaiting QML qualification as of July 2023), exhibits a computational density that is nearly two orders of magnitude smaller than that of the commercial Xilinx Zynq 7020, a popular embedded system-on-chip (SoC) architecture released in 2011 that combines a dual ARM Cortex-A9 core and an FPGA fabric, and the commercial Core i7-4610Y, a 4^{th}-generation Intel hyperthreaded dual-core processor launched in 2013 that was designed for mobile and tablet markets.

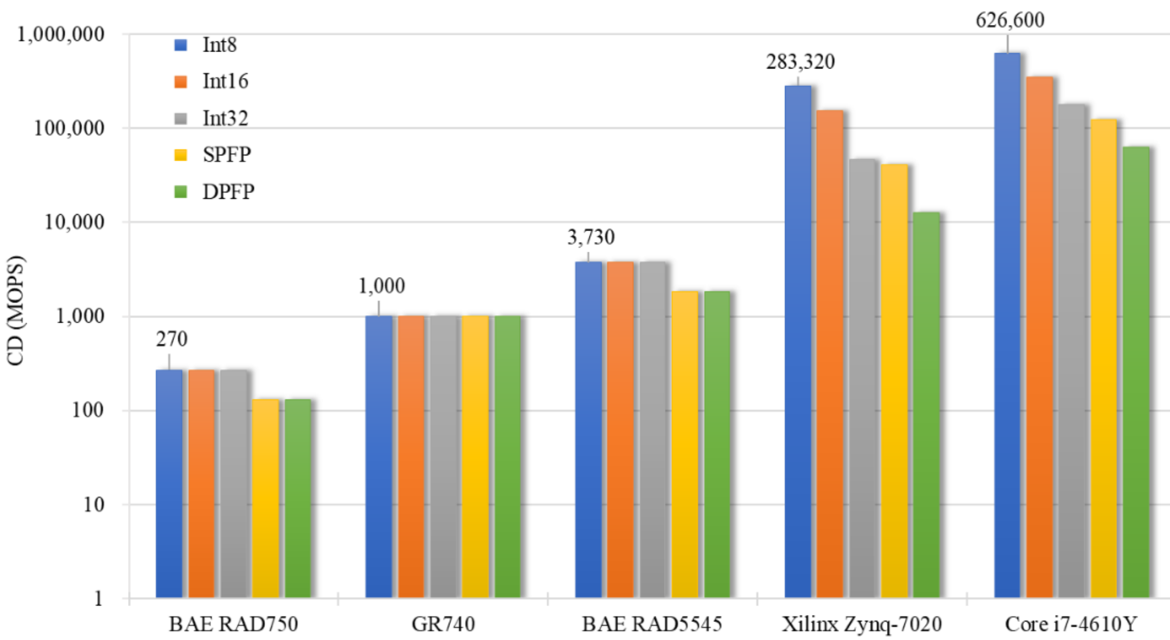

Figure 1: Comparison of Computational Density (CD) of State-of-the-art Rad-Hard Processors (BAE RAD750, CAES Gaisler GR740, and BAE RAD5545) and Commercial Embedded Processors (Xilinx Zynq 7020 and Intel Core i7-4610Y) [21]

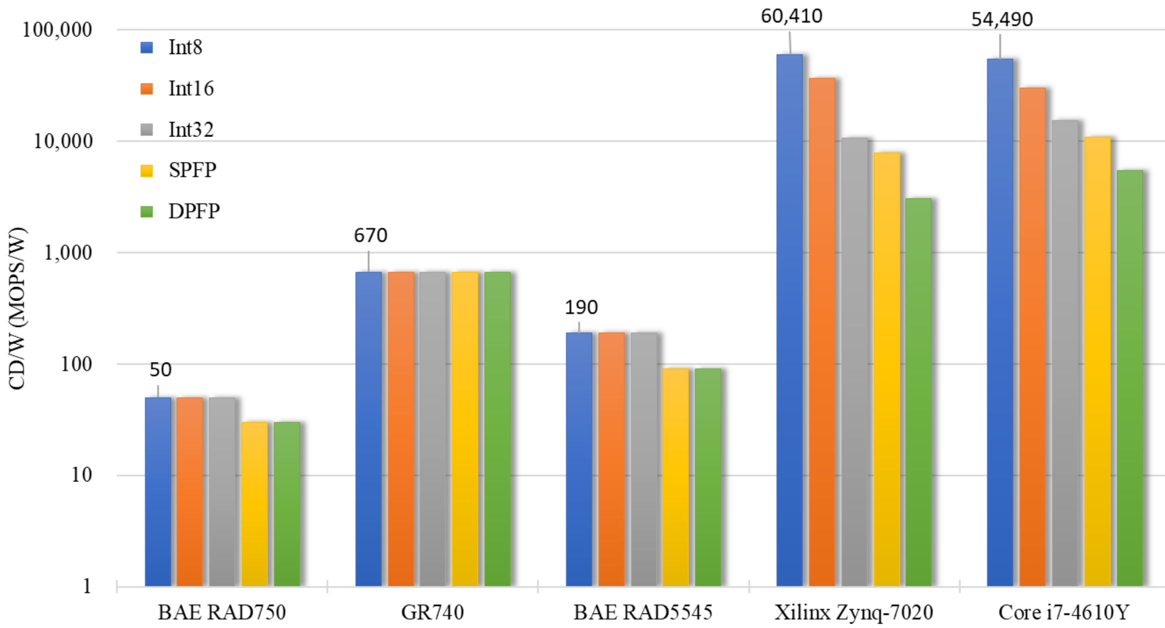

Figure 2: Comparison of Computational Density Per Watt of State-of-the-art Rad-Hard Processors (BAE RAD750, CAES Gaisler GR740, and BAE RAD5545) and Commercial Embedded Processors (Xilinx Zynq 7020 and Intel Core i7-4610Y) [21]

Likewise, the power efficiency of rad-hard processors, which can be estimated from the computational density per watt (CD/W) metric shown in Figure 2, is nearly two orders of magnitude smaller than that of commercial embedded processors [21]. Aside from the computational density and power efficiency of the processing units, memory size and bandwidth are also key considerations for AI, allowing the processor to efficiently store and stream model parameters and intermediate calculations. With rad-hard processors typically clocked at lower frequencies, their memory bandwidths also trail behind commercial DDR solutions significantly. This limited computational density and memory bandwidth severely restricts the size and complexity of AI models that can be feasibly deployed on rad-hard processors.

To reduce the disparity between rad-hard and commercial embedded technology and significantly advance the state-of-the-art in space computing, NASA and the U.S. Air Force have jointly funded the High-Performance Space Computing (HPSC) project. Originally initiated in 2013, the HPSC project has been historically plagued by significant delays. Initially contracted out to Boeing, HPSC was previously defined to include eight ARM Cortex-A53 cores arranged in two four-core clusters, known as chiplets, connected through a high-bandwidth interconnect [22]. However, with Boeing unable to deliver at key deadlines, the contract was terminated. As of August 2022, NASA has awarded the HPSC contract to Microchip, who plans to design, test, and qualify a RISC-V-based 12-core processor, comprising eight SiFive's X280 application-level cores [23]. Importantly, for fast AI model inference, the X280 cores include both SiFive Intelligence Extensions to accelerate neural network computations (touted at a performance of 4.6 trillions of

8-bit integer operations per second) and vector extensions to exploit parallelism in dataflow applications [24].

3.1 Up-screened Commercial-Off-the-Shelf Processors for Space

Due to the expense and the traditionally limited performance of rad-hard processors, low-budget, risk-tolerant Class D and small-satellite missions have gravitated toward hybrid computing designs, whereby a mix of a commercial-off-the-shelf (COTS) and rad-hard components are combined to achieve the necessary reliability and performance for the instrument or mission. The SpaceCube family of single-board computers (SBCs), developed by NASA Goddard Space Flight Center (GSFC), is a compelling example of this hybrid processing approach. Historically, this line of SBCs has employed high-performance, commercial Xilinx FPGAs and SoCs as the main processor, supplemented with peripheral rad-hard components such as voltage regulators to power the processor and watchdogs to monitor for radiation-induced upsets [25].

Of paramount importance to these hybrid processing approaches is ensuring the commercial devices are resilient to certain cumulative and acute radiation effects, which can cause erroneous bit flips or even catastrophic device failure. Thus, these commercial devices are typically up-screened for space use through accelerated, ground-based radiation testing. In terms of cumulative effects, TID testing measures the device's response to long-term ionizing damage from protons and electrons, causing parametric or functional degradation, such as transistor threshold-voltage shifts and increased device current leakage, that can lead to failure over time. In contrast, SEE testing measures the device's response to single, high-energy particle impacts, simulating heavy-ion components from cosmic rays and high-energy protons in the space environment. For CMOS-based processors, these effects can take many forms, including but not limited to:

- Single-event upsets (SEUs), which are non-destructive, soft errors that may cause bitflips in registers or memory.
- Single-event functional interrupts (SEFIs), which are soft errors that cause the device to stop functioning nominally, usually requiring a power reset to restore normal functionality.
- Single-event latchup (SEL), which results in the device operating in a high-current state that can be non-destructive (cleared by a reset) or destructive.

To ensure resiliency to various SEE-induced faults, various fault-tolerant computing techniques can be employed on these commercial devices. One common method for FPGA-based processors is triple-modular redundancy (TMR) in which the same data is sent to three identical processing circuits and the final output is selected by a majority voter.

Despite their susceptibility to radiation effects, commercial embedded devices exhibit significant improvements in performance compared to rad-hard devices as demonstrated in Figures 1 and 2, allowing for deployment of much more complex algorithms such as AI model inference. Moreover, many frameworks designed for edge AI inference already support a multitude of these commercial embedded devices. For example, the SpaceCube Mini-Z incorporates the Xilinx Zynq 7020 SoC, featuring a dual ARM Cortex-A9 CPU alongside a FPGA fabric [25]. TensorFlow Lite, a standard framework for the deployment of AI/ML applications on embedded devices, is highly optimized for the ARM Cortex-A9 architecture by specifically leveraging instructions for the Neon Single Instruction Multiple Data (SIMD) units that are designed to accelerate vector processing kernels such as matrix multiplication. Similarly, the SpaceCube v3.0 Mini SBC features the commercial Xilinx Kintex UltraScale FPGA that can instantiate the Xilinx Deep Learning Processor Unit (DPU) to accelerate neural network operations for AI/ML applications [26].

Recent research has explored the use of commercial edge AI accelerators developed for the mobile and IoT markets, including the Google Coral Edge TPU and Intel Movidius Myriad 2 and Myriad X. Because the mobile/IoT and space domains are under similar SWaP-C constraints, these edge accelerators incorporate attractive properties for space applications, demonstrating significant inferencing performance in small, low-power form factors. For example, the Google Coral Edge TPU advertises a performance of 4 trillions of 8-bit integer operations per second at only 2 W [27], with the Intel Myriad 2 and Myriad X at similar performance per watt efficiencies.

Similar to other commercial devices, these accelerators underwent radiation testing to screen them for space use, including TID and SEE (heavy ion and proton) testing. Radiation testing of the Myriad 2 was led by Ubotica and the European Space Agency [28], while efforts for the Google Edge TPU and the Myriad X were led by NASA GSFC [29]-[30]. Additionally, each of these devices has been incorporated into technology-demonstration missions to achieve flight heritage and to increase their TRL. The Myriad 2 was integrated into the Φ-sat-1 mission to perform inference of the CloudScout cloud-detection model. The SpaceCube Low Power Edge AI Resilient Node (SC-LEARN) features three Google Edge TPU in a TMR configuration and has been deployed on the Space Test Program – Houston 9 (STP-H9) pallet as part of the SpaceCube Edge Node Intelligent Collaboration (SCENIC) experiment [31]. Finally, the Myriad X has been used in multiple missions, including STP-H9/SCENIC and NASA JPL's international space station (ISS) benchmarking experiments [32].

4. **Conclusion**

Due to the unprecedented performance of deep-learning AI models in the commercial embedded domain for computer-vision and natural-language-processing tasks, space researchers are actively examining how AI can be incorporated onboard next-generation spacecraft to enhance their autonomous capabilities and science return. Unlike the commercial domain, there are a multitude of space-specific modeling and computational challenges that must be addressed. On the modeling side, large-scale datasets are typically not available for novel sensors or unexplored environments, so it can be difficult to train deep neural networks and validate their performance prior to deployment. On the computational side, traditional, rad-hard processors cannot feasibly execute standard deep-learning inference, lacking the necessary compute and memory bandwidth. High-performance, embedded COTS processors that are up-screened for space use, ranging from CPUs, GPUs, FPGAs, and custom neural-network-accelerator ASICs, offer a promising solution for more risk-tolerant missions. Unfortunately, these devices may not be able to be used in extremely harsh radiation environments. Thus, there are still significant efforts from NASA and the U.S. Air Force aimed at developing rad-hard, high-performance processors like the HPSC. Despite these challenges, the space domain is currently realizing the use of onboard AI in many technology-demonstration missions, like Φ-sat-1, in order to increase its TRL and later infuse it in higher class missions as the technology matures.

5. References

[1] Chien, S., et al. “The EO-1 Autonomous Science Agent”, *3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, New York, NY, July 18-23, 2004.

[2] Thompson, D. R., Altinok, A., Borstein, B., Chien, S. A., Doubleday, J., Bellardo, J., and K. L. Wagstaff, “Onboard Machine Learning Classification of Images by a Cubesat in Earth Orbit,” *AI Matters*, vol. 1, no. 4, pp. 38-40, Jun. 2015.

[3] Norman, C. D., et al. "Autonomous navigation performance using natural feature tracking during the OSIRIS-REx Touch-And-Go sample collection event." *The Planetary Science Journal*, vol. 3, no. 5, p. 101, 2022.

[4] Nelessen, Adam, et al. "Mars 2020 entry, descent, and landing system overview." 2019 IEEE Aerospace Conference, Big Sky, MT, Mar 2-9, 2019.

[5] “About MMS,” MMS Science Data Center. [Online]. Available: https://lasp.colorado.edu/mms/sdc/public/

[6] Giuffrida, Gianluca, et al. "The Φ-Sat-1 mission: The first on-board deep neural network demonstrator for satellite earth observation." *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1-14, 2021.

[7] Boshuizen, Christopher, et al. "Results from the planet labs flock constellation." *28th Annu. AIAA/USU Conf. on Small Satellites*, Logan, UT, Aug. 4-7, 2014.

[8] Kang, Yoojin, et al. "A deep learning model using geostationary satellite data for forest fire detection with reduced detection latency." *GIScience & Remote Sensing*, vol 59, no. 1, pp. 2019-2035, 2022.

[9] Thangavel, Kathiravan, et al. "Autonomous Satellite Wildfire Detection Using Hyperspectral Imagery and Neural Networks: A Case Study on Australian Wildfire." *Remote Sensing*, vol. 15, no.3, pp. 720, 2023.

[10] James, George L., et al. "An Efficient Wildfire Detection System for AI-Embedded Applications Using Satellite Imagery." *Fire*, vol. 6, no. 4, p. 169, 2023.

[11] Mateo-Garcia, Gonzalo, et al. "Towards global flood mapping onboard low cost satellites with machine learning." *Scientific reports*, vol. 11, no. 1, p. 7249, 2021.

[12] Syifa, Mutiara, et al. "Flood mapping using remote sensing imagery and artificial intelligence techniques: a case study in Brumadinho, Brazil." *Journal of coastal research*, 90(SI), pp. 197-204, 2019.

[13] Toupet, Olivier, et al. "A ros-based simulator for testing the enhanced autonomous navigation of the mars 2020 rover." *2020 IEEE Aerospace Conference*. Big Sky, MT, 2020.

[14] Swan, R. Michael, et al. "Ai4mars: A dataset for terrain-aware autonomous driving on mars." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[15] Chase Jr, Timothy, et al. "You Only Crash Once: Improved Object Detection for Real-Time, Sim-to-Real Hazardous Terrain Detection and Classification for Autonomous Planetary Landings." *Proceedings of AAS/AIAA Astrodynamics Specialist Conference*, 2022.

[16] "RAD6000™ Space Computers." BAE Systems. [Online]. Available: https://montcs.bloomu.edu/~bobmon/PDFs/RAD6000_Space_Computers.pdf

[17] "RAD750 Radiation-Hardened PowerPC Microprocessor." BAE Systems. [Online]. Available: https://www.baesystems.com/en-media/uploadFile/20210404045936/1434555668211.pdf

[18] "GR712RC Dual-Core LEON3-FT SPARC V8 Processor." CAES Gaisler. [Online]. Available: https://www.gaisler.com/doc/gr712rc-datasheet.pdf

[19] "GR740 Quad-Processor LEON4FT System-on-Chip Overview." CAES Gaisler. [Online]. Available: https://www.gaisler.com/doc/gr740/GR740-OVERVIEW.pdf

[20] "RAD5545™ SpaceVPX Single-Board Computer." BAE Systems. [Online]. Available: https://www.baesystems.com/en-media/uploadFile/20210404061759/1434594567983.pdf

[21] Lovelly, Tyler M., and Alan D. George. "Comparative analysis of present and future space-grade processors with device metrics." *Journal of aerospace information systems*, vol. 14, no. 3, pp. 184-197, 2017.

[22] Powell, Wesley A. "High-performance spaceflight computing (hpsc) project overview." *Radiation Hardened Electronics Technology Conference (RHET)*, no. GSFC-E-DAA-TN62651, 2018.

[23] Miller, David. "NASA Makes RISC-V the Go-to Ecosystem for Future Space Missions." SiFive, Inc. [Online]. Available: https://www.sifive.com/press/nasa-selects-sifive-and-makes-risc-v-the-go-to-ecosystem

[24] "SiFive Intelligence X280." SiFive, Inc. [Online]. Available: https://sifive.cdn.prismic.io/sifive/9405d3d0-35a1-4680-a259-7a5598d1ecb2_sifive-intelligence-x200-datasheet.pdf

[25] Brewer, Cody, et al. "NASA SpaceCube intelligent multi-purpose system for enabling remote sensing, communication, and navigation in mission architectures." *34th Annu. AIAA/USU Conf. on Small Satellites, SSC20-VI-07*, Logan, UT, Aug. 1-6, 2020.

[26] Geist, Alessandro, et al. "SpaceCube v3. 0 NASA next-generation high-performance processor for science applications." *33rd Annu. AIAA/USU Conf. on Small Satellites*, Logan, UT, Aug. 3-8, 2019.

[27] "Coral Technology: Advanced Neural Network Processing for Low-power Devices." Google Coral. [Online]. Available: https://coral.ai/technology/

[28] Buckley, Leénie, et al. "Radiation Test and in Orbit Performance of MpSoC AI Accelerator." *2022 IEEE Aerospace Conference*, Big Sky, MT, 2022.

[29] Casey, Megan, et al. "Effect of Total Ionizing Dose on Low-Power Artificial Intelligence Edge Processing Application-Specific Integrated Circuits for Space-Base Applications." *IEEE Nuclear & Space Radiation Effects Conference*, Salt Lake City, UT, 2022.

[30] Casey, Megan C., et al. "Single-Event Effects on Commercial-Off-the Shelf Edge-Processing Artificial Intelligence ASICs." *IEEE Transactions on Nuclear Science*, 2023.

[31] Goodwill, Justin, et al. "NASA spacecube edge TPU smallsat card for autonomous operations and onboard science-data analysis." *35th Annu. AIAA/USU Conf. on Small Satellites, SSC21-VII-08*, Logan, UT, 2021.

[32] Dunkel, Emily, et al. "Benchmarking deep learning inference of remote sensing imagery on the qualcomm snapdragon and intel movidius myriad x processors onboard the international space station." *IEEE International Geoscience and Remote Sensing Symposium*, 2022.