Deployable Neuromorphic AI in Space: Taxonomy and Architecture

Anonymous CVPR submission

Paper ID 29

Abstract

Neuromorphic computing is increasingly considered for onboard artificial intelligence in space due to its potential for energy-efficient, low-latency inference and its natural fit to sparse, event-driven sensing. Yet, much of the existing literature remains focused on algorithms, device characteristics, or single-application demonstrations, offering limited guidance on how neuromorphic acceleration should be integrated and operated at spacecraft system level. Consequently, architectural decisions are often technology-driven and difficult to transfer across missions and autonomy functions. This position paper addresses this gap by introducing an application-driven, system-level framework for neuromorphic on-board AI. We propose a two-dimensional taxonomy based on (i) criticality (safety-critical vs. mission-critical) and (ii) temporal behavior (persistent vs. event-driven), yielding four workload categories capturing key architectural drivers. We map these categories to reusable integration patterns, including execution modes, power management, and host–accelerator interaction assumptions. We ground the framework with feasibility-level quantitative anchors from representative workloads on neuromorphic and hybrid platforms, focusing on indicative latency, energy, and duty-cycle behavior to support early system-level architectural trade studies for autonomous spacecraft.

1. Introduction

Space missions increasingly rely on onboard autonomy under limited ground contact, tight power budgets, and growing sensing/telemetry volumes. Neuromorphic computing, which comprises spiking neural networks [22], neuromorphic hardware to run them, and event-based sensing [10], is therefore attracting interest for low-power, low-latency onboard inference, particularly for sparse or event-driven workloads [13]. In parallel, the move toward heterogeneous onboard compute stacks and AI accelerators reinforces that *system-level integration* is the key challenge, not model accuracy alone [32] [43].

A central difficulty is the diversity of autonomy functions: some are safety-relevant and run continuously or at fixed cadence, while others are mission-oriented and bursty, driven by sensing context. These differences often drive architectural choices—execution modes, host–accelerator coupling, and power management—more than the model itself. Yet much of the neuromorphic-in-space discussion remains technology- or single-use-case-driven, leaving mission designers without a consistent way to decide *when* neuromorphic acceleration is appropriate, *where* it fits in the stack, and *how* it should be operated across functions.

1.1. Related Work

Neuromorphic computing for space has been motivated in survey and perspective works that consolidate candidate mission use cases and deployment considerations [13]. In parallel, in-orbit and space-adjacent demonstrations have validated the practical relevance of event-driven sensing in operational environments, exemplified by Falcon Neuro event-based imaging experiments [25, 26]. These efforts provide strong feasibility signals for event-driven sensing pipelines, but they focus primarily on sensing, data products, and mission demonstrations rather than reusable system-engineering abstractions for integrating heterogeneous autonomy functions.

A related stream of work examines energy-efficient onboard processing and AI acceleration for satellite applications, including studies of accelerator-assisted payload processing chains [32] and neuromorphic approaches in specific onboard contexts such as communications resource management [19, 33]. In the AI4Space/CVPR workshop community, spiking neural networks have been explored for energy-efficient onboard vision tasks [18], and event-based remote sensing methods further motivate event-driven perception workloads [21]. While these works motivate feasibility and demonstrate application potential, they typically do not provide a general decision framework that connects application properties to system-level integration patterns.

1.2. Contributions

- **2D application taxonomy:** criticality × temporal behavior, yielding four workload categories (A–D).

- **Application-to-architecture mapping:** reusable integration patterns (execution modes, host coupling, power expectations) per class.
- **Reference architecture + anchors:** a modular stack and feasibility-level latency/energy/duty-cycle anchors (illustrative, not benchmarking).

2. Space Autonomy Applications

Neuromorphic and event-based techniques have found practical use across a range of space autonomy functions; the following subsections summarize the main patterns and takeaways within each function area.

Planetary landing. Scientific interest here mainly revolves around neuromorphic computing on event camera streams to either provide direct estimates of time-to-contact (TTC)[37], enabling robust navigation without full image [1], or for tackling divergence[27] , as contrast maximization on motion-compensated event frames by applying branch-and-bound methods via GPU acceleration[44].

Attitude Determination. Event cameras can act as star-tracker imagers for fast, low-power attitude sensing: events are aggregated into event frames and processed either by fused rotation estimation [4], or by progressive Hough transforms that recover relative attitude from star-streak geometry [2]. Validation uses real starlight with Earth rotation as a precise reference achieving about 18 arcsec RMSE. [28].

Pose Estimation. Neuromorphic computing targets directly rendezvous and on-orbit servicing use cases using event cameras to overcome classical frame camera artifacts. Recent work either uses reconstructed event frames to run conventional neural pose estimation algorithms [5][35], or converts events into intensity-like images to reuse classic structure-from-motion for trajectory and 3D reconstruction [23]. Domain-gap studies to reduce sim-to-real fragility produced the SEENIC [14], FRESH [15] and SPADES [35] datasets, among the most relevant public resources.

On-Board SSA. On-board space situational awareness for nano-satellites increasingly uses event cameras to autonomously detect and characterize debris while limiting power usage and downlink; EventSat demonstrate on-board inference in LEO despite highly dynamic imagery [7][36]. Payload studies quantify detectability via sensor characteristics and size baffles/vanes for stray-light suppression, while COTS neuromorphic experiments confirm viability but expose low-light noise and low-dynamics degradation[11].

EO Payloads. Event cameras have already flown for lightning and other fast atmospheric onsets, where microsecond latency and sparse output cut exposure and downlink[3]. ISS lightning demos further support neuromorphic sensing as a feasible EO instrument class [26]. For radiometric scenes, events are fused with images to reconstruct single-exposure HDR products while addressing image–event domain gaps and sparsity[21]. System concepts also use events as always-on triggers for change detection concepts[38], while onboard research spans SNN land-cover classification[17], physics-aware multispectral disaster change mapping to prioritize downlink[45], and neuromorphic processor architectures for methane payloads[29].

Communications. Neuromorphic SatCom targets on-board RRM, beamforming, and spectrum awareness where classical optimization or dense CNNs are too costly: Spiking RRM models on Loihi-2 show competitive accuracy at far lower power than CNN baselines [33]. Sparse spiking beamforming maps LASSO-style designs onto locally competitive networks to reduce active RF chains with minimal performance loss [24] while interference is cast as neuromorphic classification, with large power savings on reduced-DFT features[8]. Enabling work defines on-board processor KPIs[32], explores split SNN inference over optical ISL to reduce inter-satellite transmission[39], and benchmarks neuromorphic processors against conventional AI accelerators[19] and datasets [42].

FDIR. Neuromorphic FDIR targets always-on, low-SWaP on-board anomaly detection by replacing threshold logic with learned nominal-behaviour models[31]. A common approach is reconstruction-based autoencoders, where anomalies appear as elevated reconstruction error; spiking variants leverage sparse spike-coded representations for streaming inference under tight compute and memory [40]. Flight relevance is strengthened by benchmarking ML-enabled FDIR on low-power boards and by FPGA mappings of spiking autoencoders reporting utilization and power compatible with continuous on-board execution [30][34].

3. Taxonomy and Application-to-Architecture Mapping for Neuromorphic AI

While all works presented in Section 2 highlight feasibility and promise, they do not provide a compact system-level abstraction that helps mission designers reason about *how* heterogeneous autonomy applications should be integrated and operated within a spacecraft compute stack. In this section, we introduce a two-dimensional taxonomy that captures two dominant architectural drivers across autonomy functions: **criticality** and **temporal behavior**. We then use this taxonomy as the basis for a structured application-to-architecture mapping.

3.1. Criticality × Temporal Behavior

The proposed taxonomy classifies onboard autonomy applications along two axes: (i) **criticality**, distinguishing functions that are safety-critical (spacecraft protection and fault containment) from those that are mission-critical (per-

Table 1. Application-to-architecture mapping for representative space autonomy functions. Hybrid strategy denotes integration patterns (triggering/gating, voting/diversity, distillation). References for each row: Landing/TRN [27, 37, 44], Star tracking [2, 20, 28], SSA [25, 36], EO [18, 21], Comms [19, 32, 33], FDIR [13]. Categories A–D reflect criticality and temporal behavior and guide expected execution mode, power profile, and host coupling at the system level.

Function	Cat	Mode	Power	Host	Hybrid
Landing/TRN ; Pose/VBPE [27, 37, 44]	D	Burst	Idle≈0, bursts	Loose–mod	Triggering; Distillation
Star tracking [2, 20, 28]	C	Cont.	Steady low	Mod–tight	Vote; Distillation
SSA (windowed obs.) [25, 36]	D	Windowed	Low idle, bursts	Loose	Triggering; Vote
EO payloads [18, 21]	C/D	Cont./burst	Steady/duty	Loose	Distillation; Triggering
Comms (monitor/adapt) [19, 32, 33]	C/D	Periodic+evt	Steady+spikes	Tight	Trig; Vote/Distillation
FDIR (monitor/safe) [13]	A/B	Periodic+evt	Steady / idle≈0	Tight	Vote; Triggering

formance, science return, or operational efficiency), and (ii) **temporal behavior**, distinguishing **persistent/periodic** processing from **event-driven/sporadic** processing. The second axis is particularly relevant for neuromorphic systems because event-driven sensing and sparse computation enable activity-proportional resource usage [13, 25].

We define the **Categories** as:

A. **Safety-critical & persistent:** continuous or periodic monitoring and protection functions (e.g., health monitoring / protection loops), where predictable execution and conservative behavior dominate.

B. **Safety-critical & event-driven:** contingency-triggered safeguards that remain near-idle until a detected condition requires fast response (e.g., event-triggered fault containment).

C. **Mission-critical & persistent:** sustained onboard processing pipelines supporting mission objectives under continuous resource budgets (e.g., persistent payload intelligence or comms intelligence), often benefiting from energy-efficient acceleration [32, 33].

D. **Mission-critical & event-driven:** bursty, context-driven perception or data-reduction workloads, commonly aligned with event-based sensing and sparse inference (e.g., event-driven perception and onboard data triage) [18, 21, 25].

3.2. From Taxonomy to Architecture Mapping

The role of this taxonomy is not to categorize missions exhaustively, but to provide a *decision lens* that links application properties to architectural consequences. In the next subsection, we map Categories A–D to system-level integration patterns (execution modes, power profiles, host coupling, and optional hybrid strategies), enabling consistent architectural trade-offs for neuromorphic onboard AI across heterogeneous autonomy functions [19, 32].

3.3. Application-to-Architecture Mapping

Table 1 operationalizes the taxonomy by mapping applications to criticality x temporal-behavior categories (A-D). It highlights execution mode, power profile and how neuromorphic is coupled to the host compute. It also indicates a proposed hybrid collaboration strategy: neuromorphic triggers the host (*triggering*), the host distills knowledge to neuromorphic (*distillation*), or a voting scheme combines both (*voting*).

4. Reference Architecture, Feasibility Anchors

We introduce a layered reference architecture for neuromorphic onboard AI and summarize feasibility anchors for latency, energy, and scaling, providing a compact systems view for early mission design.

4.1. From Classical Stacks to Neuromorphic

Modern spacecrafts employ heterogeneous stacks of sensors, processors, accelerators, and supervisory flight computers [32] [16]. Neuromorphic computing extends this model with an execution paradigm optimized for sparse, event-driven workloads [6, 12]. Rather than replacing conventional pipelines, neuromorphic processors complement existing accelerators by enabling activity-proportional inference within heterogeneous onboard AI systems.

4.2. A Layered Neuromorphic Spacecraft Stack

We propose a layered neuromorphic stack that preserves sparsity across the sensing-to-decision pipeline while remaining compatible with existing architectures. The model emphasizes sparsity preservation, asynchronous execution, and duty-cycle-aware orchestration.

Layer 1: Sensing. Event-driven sparsity originates from native event sensors or eventized dense signals. Event cameras have demonstrated in-orbit feasibility [1, 25, 37].

Layer 2: Front-end Encoding. Preprocessing regulates activity via filtering or temporal aggregation, analogous to FPGA-style front ends [27, 44].

Layer 3: Neuromorphic Acceleration. Asynchronous inference is performed using spiking or event-driven models, exploiting sparse activation for energy efficiency [6, 9].

Layer 4: Host Supervision. A supervisory processor manages orchestration, redundancy, and interfaces, enabling hybrid neuromorphic–classical pipelines [32, 41].

Layer 5: Autonomy Integration. Outputs feed GNC, perception, and FDIR loops, typically augmenting rather than replacing classical decision stacks [13].

Cross-layer insight. Neuromorphic components integrate into hybrid stacks rather than operating in isolation. Recurring patterns include hierarchical triggering, redundancy across paradigms, and gradual migration from deep learning to spiking implementations. Such hybridization combines low-latency perception with the robustness of conventional accelerators.

4.3. Operating Regimes

While Section 3 establishes the taxonomy and mapping, the question is how these categories manifest in practice. Fig. 1 provides empirical anchors by situating representative operating regimes within a latency–power design space.

The power–latency map reveals continuous operating envelopes rather than discrete clusters. Persistent workloads occupy stable mid-latency, low-power regions consistent with steady-state neuromorphic co-processing [6, 9], while event-driven pipelines span broader vertical envelopes with low idle power and brief latency excursions, reflecting activity-proportional execution [19, 37]. Sensing-dominated front ends define the lower-latency frontier, where sparse acquisition gates heavier downstream processing [44]. These regimes align with activity-dependent scaling in neuromorphic systems: latency grows with temporal integration depth, energy with spike counts, and power follows duty cycle [12].

From a design perspective, temporal workload behavior emerges as a strong indicator of operating regime: event-driven workloads favor low-duty-cycle, bursty operation, while persistent workloads align with steadier power–latency bands. This trend is visible in Fig. 1, where event-driven examples (e.g., points 1–5 and 9–12) span broader power ranges, while more persistent workloads (e.g., points 6–8) cluster into tighter bands

5. Discussion: Design Implications

End-to-end sparsity preservation. A primary system-level takeaway is that neuromorphic advantage is rarely delivered by a single component in isolation. Instead, efficiency emerges when sparsity is preserved across the full pipeline (sensing/encoding → inference → decision). Converting sparse signals into dense intermediate representations too early can shift the dominant cost to data movement and host-side preprocessing, eroding the benefits implied by Fig. 1 and the integration patterns in Table 1.

Criticality drives coupling and assurance. The taxonomy highlights that safety-critical categories (A/B) require tighter coupling to supervisory control, conservative confirmation paths, and explicit fallback behaviors. In contrast, mission-critical categories (C/D) can exploit duty-cycling and context-driven activation, and can tolerate graceful degradation (e.g., reduced fidelity or reduced duty cycle) without compromising spacecraft safety.

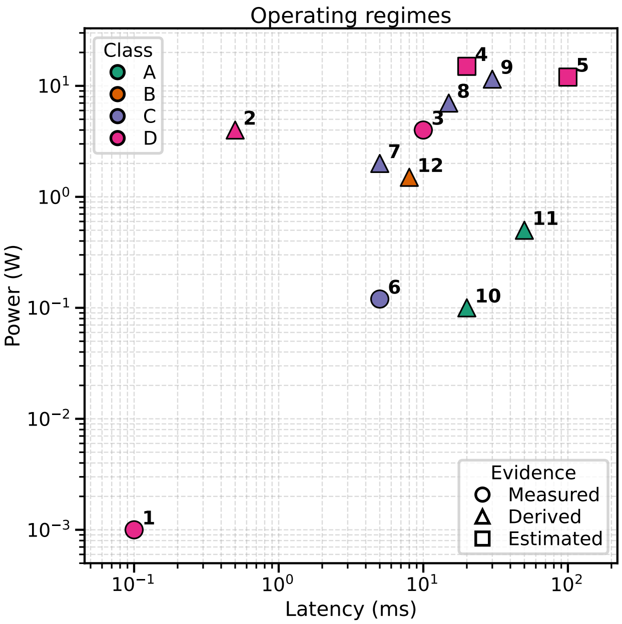

Figure 1. Neuromorphic spacecraft operating regimes (log–log latency vs. power). Numbers correspond to representative systems: (1) ESA optic-flow landing [44]; (2–3) Falcon Neuro ISS sensing and lightning observations [25, 26]; (4) EvSat3D event-based pose estimation [23]; (5) event-based divergence landing [27]; (6) Loihi-based onboard RRM [33]; (7) Brainsat neuromorphic onboard computer [29]; (8) event-based nanosat pointing [20]; (9) EventSat payload studies [7]; (10–11) neuromorphic onboard monitoring and FDIR architectures [30, 31]; (12) event-driven nanosatellite safeguard / mission-architecture anchor [36]. Marker shape denotes evidence type (measured/derived/estimated), and color indicates functional category.

Hybrid strategies are integration tools. Triggering/gating, voting/diversity, and distillation are best viewed as system-integration patterns that improve robustness or fit within tight compute/energy envelopes (Table 1), rather than as new algorithms. These patterns help reconcile heterogeneous autonomy requirements with practical constraints on embedded hardware and operations.

When neuromorphic is not the right fit. Sustained, dense pipelines with limited intrinsic sparsity may be better served by conventional accelerators unless eventization/encoding preserves meaningful sparsity. The proposed framework is intended to make such cases explicit early, supporting technology selection rather than assuming neuromorphic is universally optimal.

References

[1] Loïc J. Azzalini, Emmanuel Blazquez, Alexander Hadjiivanov, Gabriele Meoni, and Dario Izzo. On the generation of a synthetic event-based vision dataset for navigation and landing, 2023. 2, 3

[2] Samya Bagchi and Tat-Jun Chin. Event-based star tracking via multiresolution progressive hough transforms, 2019. 2, 3

[3] Olivier Chanrion, Nicolas Pedersen, and Torsten Neubert. Cam-davis executive summary report. Executive Summary Report CAMDAVIS-DTU-ESR-01, DTU Space, Technical University of Denmark, 2024. ESA Contract No. 4000139066/22/NL/GLC/ov. 2

[4] Tat-Jun Chin, Samya Bagchi, Anders Eriksson, and Andre van Schaik. Star tracking using an event camera, 2019. 2

[5] Jonathan Courtois, Benoît Miramond, and Alain Pegatoquet. Spiking monocular event based 6d pose estimation for space application. 2024. 2

[6] Mike Davies et al. Loihi: A neuromorphic manycore processor. *IEEE Micro*, 2018. 3, 4

[7] Sydney Dolan, Lara Schuberth, Rugved Arge, Ramón María García Alarcia, Vincenzo Messina, Clemente J. Juan Oliver, Francesco Salmaso, Jaspar Sindermann, Federico Sofio Avogadro, and Alessandro Golkar. Design and analysis of an event camera payload for space-based object detection on the eventsat 6u cubesat mission. In *15th IAA Symposium for Small Satellites for Earth System Observation*, Berlin, Germany, 2025. 2, 4

[8] Geoffrey Eappen, Saed Daoud, Nicolas Skatchkovsky, Flor Ortíz, Eva Lagunas, Wallace Alves Martins, and Symeon Chatzinotas. Neuromorphic models for energy-efficient onboard interference detection in satellite systems. *IEEE Transactions on Aerospace and Electronic Systems*, 61: 17682–17702, 2025. 2

[9] Charlotte Frenkel et al. A 0.086-mm² 12.7-pj/sop 64k-synapse neuromorphic processor. *Nature Electronics*, 2019. 3, 4

[10] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jorg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP:1–1, 2020. 1

[11] Jessica L. Horn. Neuromorphic vision sensors for space-based applications. Master of science thesis, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio, 2021. AFIT-ENY-MS-21-M-304. 2

[12] Giacomo Indiveri and Shih-Chii Liu. Memory and information processing in neuromorphic systems. *Proceedings of the IEEE*, 2015. 3, 4

[13] D. Izzo et al. Neuromorphic computing and sensing in space, 2022. 1, 3, 4

[14] Mohsi Jawaid, Ethan Elms, Yasir Latif, and Tat-Jun Chin. Towards bridging the space domain gap for satellite pose estimation using event sensing, 2022. 2

[15] Mohsi Jawaid, Ethan Elms, Yasir Latif, and Tat-Jun Chin. Towards bridging the space domain gap for satellite pose estimation using event sensing, 2022. 2

[16] Antonis Karteris, Alexis Chatzistylianos, Evgenios Tsigkanos, Giannis Panagiotopoulos, Alexandros Stavropoulos, Mathieu Bernou, Dimitrios Soudris, and George Lentaris. Satellite-as-a-service architecture for ml edge computing on heterogeneous processing platforms. In *2025 European Data Handling Data Processing Conference (EDHPC)*, pages 1–5, 2025. 3

[17] Andrzej Kucik and Gabriele Meoni. Investigating spiking neural networks for energy-efficient on-board ai applications. a case study in land cover and land use classification. 2021. 2

[18] A. S. Kucik and G. Meoni. Investigating spiking neural networks for energy-efficient on-board ai applications: A case study in land cover and land use classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021. 1, 3

[19] E. Lagunas et al. Performance evaluation of neuromorphic hardware for onboard satellite communication applications. *IEEE Wireless Communications*, 2024. 1, 2, 3, 4

[20] Yasir Latif, Peter Anastasiou, Yonhon Ng, Zebb Prime, Tien-Fu Lu, Matthew Tetlow, Robert Mahony, and Tat-Jun Chin. High frequency, high accuracy pointing onboard nanosats using neuromorphic event sensing and piezoelectric actuation, 2023. 3, 4

[21] X. Li et al. Ers-hdri: Event-based remote sensing hdr imaging. *Remote Sensing*, 2024. 1, 2, 3

[22] Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10 (9):1659–1671, 1997. 1

[23] Shahid Shabeer Malik, Maryam Moshrefizadeh, Omar Tahri, Xiaoli Bai, Erik Blasch, Vasit Sagan, and Hadi AliAkbarpour. Evsat3d: Satellite pose estimation and 3d reconstruction with event camera. *IEEE Access*, 13:130340–130352, 2025. 2, 4

[24] Wallace Martins, Eva Lagunas, Nicolas Skatchkovsky, Flor Ortiz, Geoffrey Eappen, Osvaldo Simeone, Bipin Rajendran, and Symeon Chatzinotas. Satellite adaptive onboard beamforming using neuromorphic processors. pages 1–6, 2024. 2

[25] M. G. McHarg et al. Falcon neuro: an event-based sensor on the international space station. *Optical Engineering*, 2022. 1, 3, 4

[26] M. G. McHarg et al. Falcon neuro space-based observations of lightning using event-based sensors. *Frontiers in Remote Sensing*, 2024. 1, 2, 4

[27] Sofia McLeod, Gabriele Meoni, Dario Izzo, Anne Mergy, Daqi Liu, Yasir Latif, Ian Reid, and Tat-Jun Chin. Globally optimal event-based divergence estimation for ventral landing, 2022. 2, 3, 4

[28] Dennis Melamed, Connor Hashemi, and Scott McCloskey. Quantifying accuracy of an event-based star tracker via earth's rotation, 2025. 2, 3

[29] Raphael Mena Morales, Pablo Miralles, Diviya Devani, Eric Knappenberger, Michael Skreta, Georgios Metaxas, and Prerna Baranwal. Brainsat: Hardware development of a neuromorphic on-board computer applied to methane detection from low earth orbit. In *Proceedings of the 75th International Astronautical Congress (IAC), IAF Space Systems*

Symposium (D1), pages 1198–1212, Milan, Italy, 2024. International Astronautical Federation. 2, 4

[30] James Murphy, John E. Ward, and Brian Mac Namee. Low-power boards enabling ml-based approaches to fdir in space-based applications. In *Proceedings of the 35th Annual Small Satellite Conference*, Logan, Utah, USA, 2021. SSC21-S1-46. 2, 4

[31] NASA. Neuromorphic spacecraft fault monitor. NASA TechPort Project 102930, 2023. STTR Phase II, Award ID 80NSSC23CA126. 2, 4

[32] F. Ortiz et al. Onboard processing in satellite communications using ai accelerators. *Aerospace*, 2023. 1, 2, 3

[33] F. Ortiz et al. Energy-efficient on-board radio resource management for satellite communications via neuromorphic computing. *IEEE Transactions on Machine Learning in Communications and Networking*, 2024. 1, 2, 3, 4

[34] D. Conficconi P. Ritirato, G. Sorrentino. Spiking autoencoder architecture for anomaly detection on satellite data with fpgas. In *SpacE FPGA Users Workshop (SEFUW), 6th Edition*, 2025. Contribution in the Artificial Intelligence/Machine Learning track, ESA SEFUW 6th Edition. 2

[35] Arunkumar Rathinam, Haytam Qadadri, and Djamila Aouada. Spades: A realistic spacecraft pose estimation dataset using event sensing, 2023. 2

[36] Lara Schuberth, Vincenzo Messina, Ramón María García Alarcia, Jaspar Sindermann, Kian Bostani Nezhad, Roberto Aldea Velayos, Sofia Lennerstrand, Julie Rollet, Federico Sofio, Alessandro Tinucci Monibas, Leonhard Kessler, and Alessandro Golkar. Leveraging event-based cameras for enhanced space situational awareness: A nanosatellite mission architecture study. In *Proceedings of the 75th International Astronautical Congress (IAC)*, Milan, Italy, 2024. International Astronautical Federation. Paper ID: IAC-24,A6,IP,54,x84856. 2, 3, 4

[37] Olaf Sikorski, Dario Izzo, and Gabriele Meoni. Event-based spacecraft landing using time-to-contact. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1941–1950, 2021. 2, 3, 4

[38] Stephen Smith, Cormac Purcell, and Zdenka Kuncic. Training-free ai for earth observation change detection using physics aware neuromorphic networks. *Scientific Reports*, 15(1), 2025. 2

[39] Zihang Song and Petar Popovski. Neuromorphic split computing via optical inter-satellite links, 2025. 2

[40] Peter Stratton, Andrew Wabnitz, and Tara Hamilton. A spiking neural network based auto-encoder for anomaly detection in streaming data. pages 1981–1988, 2020. 2

[41] Vivienne Sze et al. *Efficient Processing of Deep Neural Networks*. Morgan Claypool, 2020. 3

[42] Evgenios Tsigkanos, Giannis Panagiotopoulos, Giorgos Ioannopoulos, Romain Clement de Givry, Petra Jory, Jakub Nalepa, and Tomas Navarro Reverte. Synthetic data generators for enhanced space-based network traffic modeling. In *2025 European Data Handling Data Processing Conference (EDHPC)*, pages 1–7, 2025. 2

[43] Evgenios Tsigkanos, Giannis Panagiotopoulos, Alexander Klaser, Adrian Leu, Mathieu Bernou, and Christos Tsigkanos. Towards verification of autonomous machine learning systems on-board spacecraft. In *2025 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C)*, pages 205–206, 2025. 1

[44] F. Valette, F. Ruffier, and S. Viollet. Neuromorphic Computation of Optic Flow Data. techreport 08/6303b, European Space Agency, the Advanced Concepts Team, 2009. Available online at www.esa.int/act. 2, 3, 4

[45] Diego Wildenstein et al. Stp-h12-vantage: Combining visual and event-based sensing for earth observation. In *Proceedings of the Small Satellite Conference*, Salt Lake City, UT, USA, 2025. Session VII: Science Mission Payloads – Enterprise. 2