Autonomous Perception and Onboard Intelligence for Space Missions: A Survey

Anonymous CVPR submission

Paper ID 37

Abstract

The integration of deep learning into space missions is revolutionizing autonomous operations, from planetary exploration to on-orbit servicing and Earth observation. However, deploying computer vision systems in space presents unique challenges, including stringent Size, Weight, Power, and Cost (SWaP-C) constraints, extreme illumination variations, and a profound "sim-to-real" domain gap due to the scarcity of in-orbit data. This paper provides a comprehensive review of the rapidly evolving landscape of autonomous space perception, structured around five core pillars: spacecraft pose estimation, multi-modal sensing, onboard edge computing, vision-based navigation, and mission robustness. We analyze recent advancements in lightweight architectures, event-based perception, and hardware-aware optimizations that enable real-time inference on radiation-tolerant edge accelerators. Furthermore, we examine the critical role of high-fidelity synthetic data generation and domain adaptation techniques in bridging the reality gap. Finally, we discuss open challenges such as adversarial robustness and uncertainty quantification, proposing a phased research roadmap toward collaborative and certifiable in-orbit AI ecosystems.

1. Introduction and Scope

The integration of deep learning (DL) into space missions is driven by the demand for advanced autonomy in planetary exploration, proximity operations, and Space Domain Awareness (SDA) [1, 6]. Applications such as Terrain Relative Navigation (TRN) for lunar and Martian landings [6], and autonomous rendezvous and docking (ARD) for On-Orbit Servicing (OOS) [39, 40], require robust computer vision systems capable of operating under extreme constraints. Unlike terrestrial environments, spaceborne systems must adhere to stringent Size, Weight, Power, and Cost (SWaP-C) limitations [6], often relying on radiation-hardened electronics that lag significantly behind commercial GPUs in performance [6].

The operational landscape is further complicated by harsh environmental factors, including extreme illumination variations, deep shadows, and the absence of atmospheric diffusion [1, 40, 47]. High orbital velocities and tumbling motions of non-cooperative targets, such as space debris or defunct satellites, necessitate high-frequency perception updates (10–25+ Hz) and resilient feature tracking, often favoring edge-based detection over traditional interest points [33, 44, 62]. Over the past decade, AI for space has evolved from foundational CNN adaptations for pose estimation [1, 62] to the deployment of advanced architectures like Transformers, Deep Reinforcement Learning (DRL) for autonomous scheduling, and Vision-Language Models (VLMs) for human-in-the-loop interaction [45, 62, 65] (see Fig. 1). Despite these milestones, significant hurdles remain, including the "sim-to-real" domain gap caused by the scarcity of high-quality orbital datasets [1, 62] and the limited set of operations supported by onboard accelerators like the Edge TPU [6].

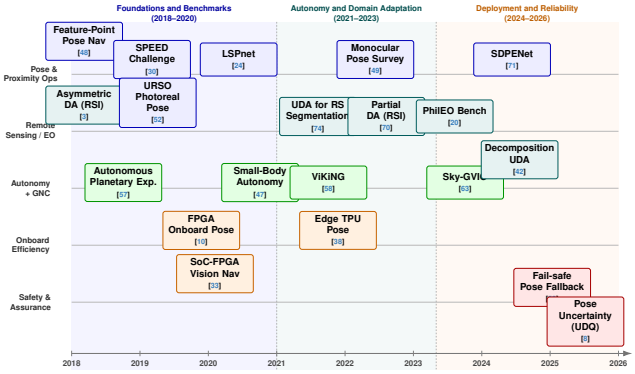

Figure 1. Chronological trajectory of space-based AI research (2018–2026).

To provide a structured overview of this rapidly evolving field, we categorize the landscape of space-based autonomous systems into five core pillars: spacecraft pose estimation, remote sensing, autonomous navigation (GNC), onboard efficiency, and mission assurance (see Fig. 2). This taxonomy serves as the organizational framework for the subsequent sections of this review.

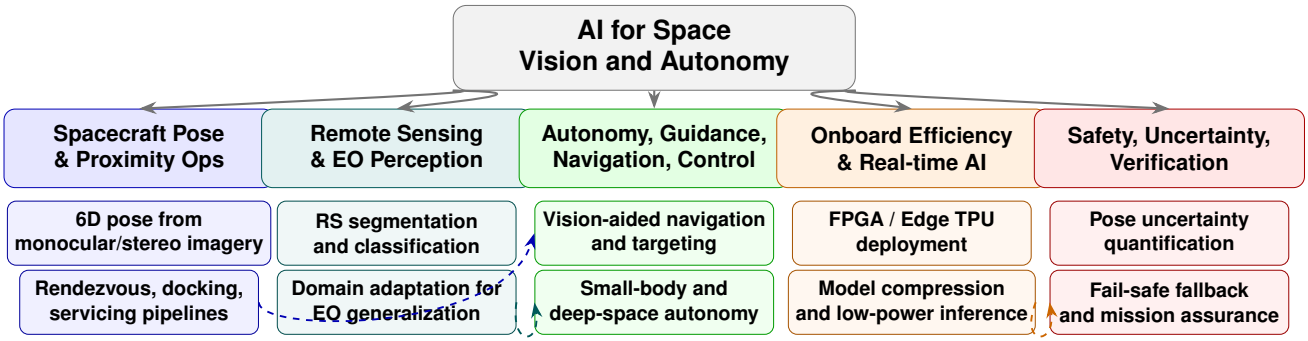

Figure 2. Task-centric taxonomy for autonomous space perception literature. The structure highlights how core mission problems connect to deployment constraints and operational assurance.

2. Spacecraft Pose Estimation and Methods

Precise 6-DoF localization underpins autonomous proximity operations, balancing geometric rigor with computational efficiency.

2.1. Hybrid Pipelines vs. End-to-End Regression

Spacecraft pose estimation architectures generally follow two paradigms: hybrid modular pipelines and direct end-to-end regression [60, 68, 72]. Hybrid approaches decompose 6-DoF estimation into sequential stages, typically involving object detection for Region of Interest (RoI) identification, 2D semantic keypoint regression, and a geometric Perspective-n-Point (PnP) solver [50, 60]. These pipelines combine the feature extraction capabilities of deep learning with the geometric rigor of solvers like Efficient PnP (EPnP) [50, 62]. Conversely, direct regression methods map input images to pose parameters using a single neural network, bypassing the need for 3D CAD models or manual feature engineering [60, 62]. While hybrid methods offer superior interpretability and accuracy by isolating errors within modules [24, 62], direct methods like LSPnet provide the computational efficiency required for resource-constrained embedded systems [24, 72]. Recent architectures, such as SDPENetv2, have significantly narrowed this accuracy gap through advanced feature utilization [73].

2.2. Keypoint Detection and Landmark Regression

Landmark localization is fundamental to hybrid pipelines, often employing high-resolution architectures like HRNet to preserve spatial precision [7, 28, 30]. Keypoints are typically represented as 2D Gaussian heatmaps and processed via PnP solvers [28, 30]. On the SPEED dataset, HRNet-based approaches have achieved rotational errors as low as 0.41° and translational errors of 0.032 m [30]. To balance performance and efficiency, lightweight models such as tiny-YOLOv3 achieve real-time inference (50 FPS) with minimal parameter counts (0.89M) [28]. Bridging the synthetic-to-real domain gap remains a primary focus, with techniques including self-supervised pseudo-labeling, adversarial training, and specialized augmentations like CropAndPad [35, 62, 72]. Nonlinear refinement, such as Simulated Annealing Levenberg-Marquardt (SA-LMPE), is often essential for achieving competitive accuracy in international challenges [7].

2.3. Orientation and Attitude Parameterization

Attitude estimation is complicated by the non-Euclidean nature of rotation spaces. Quaternions are common for direct regression but suffer from non-injectivity and slow convergence; thus, cosine-based angular distance losses are often preferred [52, 73]. Soft classification techniques, which map orientations to discrete probability mass functions, allow for multi-hypothesis extraction in ambiguous scenarios [52]. The Spacecraft Discrete Pose Estimation Network (SDPENet) represents Euler angles as discrete distributions, avoiding singularities and gimbal lock while reducing parameter counts [72, 73]. Its successor, SDPENetv2, incorporates a learnable token head and discrete pose weights in spherical coordinates, achieving state-of-the-art results on the SPEED benchmark [73].

2.4. Temporal and Sequential Pose Estimation

Exploiting temporal continuity is vital for robust navigation during rendezvous and proximity operations. Architectures like ChiNet utilize Recurrent Convolutional Neural Networks (RCNNs) with LSTM units to model long-term dependencies in multi-modal sequences [68]. Motion-aware frameworks integrate Vision Transformers (ViT) with optical flow to enforce temporal consistency, outperforming static single-image baselines [60]. Recent challenges, such as SPARK 2024, have demonstrated the utility of motion consistency for self-iterative optimization and pseudo-label smoothing [68]. The Bennunet architecture further integrates time-distributed CNNs and Gated Recurrent Units (GRUs) to extract velocity and rotation, providing substantial gains in highly constrained scenarios [37].

2.5. Transformer and Attention Mechanisms

Transformer-based architectures have emerged as powerful alternatives to CNNs, offering superior long-range dependency modeling and cross-domain generalization [60, 62]. Vision Transformers (ViTs) are increasingly used as backbones, often fused with motion cues for temporal correlation [60]. Specialized models like SU-Net integrate Multi-Head Self-Attention Blocks (MHSAB) into dense residual U-Nets to extract global information from sparse ISAR imagery [22]. Attention mechanisms, including Spatial-Semantic Interaction Attention (SSIA) and Coordinate Attention, refine feature fusion in lightweight networks without significant overhead [72]. These models have demonstrated competitive performance on standard CPUs, making them suitable for onboard edge computing where GPU acceleration is unavailable [25].

3. Synthetic Data and Domain Gap

Bridging the 'sim-to-real' divide remains the primary bottleneck for deploying robust vision models in the data-scarce environment of deep space.

3.1. Synthetic Dataset Generation and Rendering

The scarcity of real-world orbital imagery has driven the development of high-fidelity simulators to generate synthetic training data. URSO (Unreal Rendered Spacecraft on Orbit) utilizes Unreal Engine 4 to model Low Earth Orbit (LEO) conditions, employing physically based rendering (PBR) and space-specific lighting, including directional sunlight and Earth albedo, while disabling ambient light to replicate high-contrast orbital environments [52]. The SurRender simulator by Airbus employs ray-tracing and energy-conservative Bidirectional Reflectance Distribution Functions (BRDFs) to model complex reflections and planetary surfaces via multi-resolution Digital Elevation Models (DEMs) [32, 47]. For proximity operations near minor bodies, SPyRender integrates with the SPICE Toolkit to compute precise object states, using PBR materials to replicate the radiometric properties of asteroids and comets [36, 37].

Hardware-in-the-Loop (HIL) facilities, such as DLR's TRON and the Airbus Robotic Laboratory, bridge the reality gap using physical mockups, robotic arms, and specialized illumination (e.g., xenon lamps for solar radiation) to capture realistic imagery with ground-truth established via motion capture [30, 32, 40]. Advanced frameworks like RAPTOR and SPICaM further enhance realism by modeling sensor-level degradations, including dark current, readout noise, and Photo-Response Non-Uniformity (PRNU), alongside physically based camera models that convert light flux into electronic signals [21, 47]. These HIL-generated datasets are critical for training models that generalize to real mission data, such as Chang'e-3 and LROC imagery [32].

3.2. Benchmark Datasets

Standardized benchmarks are pivotal for quantifying the synthetic-to-real domain gap (see Fig. 3). The Spacecraft Pose Estimation Dataset (SPEED) provided 15,000 synthetic images and 300 real images captured at the TRON facility, focusing on the Tango spacecraft [30, 40]. Its successor, SPEED+, expanded this to three domains: synthetic OpenGL renderings, Lightbox (HIL with diffuse lighting), and Sunlamp (HIL with high-contrast arc lamps), specifically designed to evaluate robustness across varying illumination [21, 35]. The SPARK 2024 challenge pushed the field further by requiring cross-domain generalization from Unity3D simulations to real laboratory data, where top-performing methods utilized style augmentation and self-supervised adaptation to achieve nearly tenfold accuracy improvements over direct regression baselines [35, 68].

Beyond RGB sensors, the SPADES (Spacecraft Pose Estimation Dataset using Event Sensing) dataset introduces asynchronous event-based data, offering high temporal resolution and wide dynamic range (HDR) [54]. SPADES proposes a 3-Channel (3C) event representation to leverage established 2D CNN architectures, demonstrating that event-based sensing can narrow the Sim2Real gap by reducing sensitivity to extreme lighting [54, 60]. Evaluation typically relies on the SLAB score, a composite of normalized translation and rotation errors, or the Median Normalized Pose Error (MNPE) to mitigate outliers in complex orbital scenes [30, 50].

3.3. Domain Adaptation and Sim-to-Real Transfer

Bridging the Sim2Real gap is critical for deploying models trained on synthetic data to in-orbit scenarios. Unsupervised Domain Adaptation (UDA) techniques, such as adversarial training using Mask-CycleGAN, align synthetic and real data distributions, though they often incur high computational costs [35, 54, 62]. Self-supervised frameworks like UF-SPE integrate offline domain generalization via style augmentation with online UDA using RANSAC EPnP to generate pseudo-labels on unlabeled real imagery [35]. This approach significantly reduces attitude and position errors on benchmarks like SPEED+ while maintaining computational efficiency suitable for in-orbit execution [35].

Other strategies focus on extracting domain-invariant features through knowledge distillation or Output-Guided Biased Sampling (OGBS) to enhance cross-domain robustness [62]. For event-based perception, architectures like MV-DA4E utilize Multi-View Pooling and Domain Adaptation Blocks (DABlocks) incorporating Gradient Reversal Layers (GRL) and Maximum Mean Discrepancy (MMD) to align feature distributions between simulated and real

Dataset / Benchmark	Primary Task	Modality	Realism	Annotation Depth	Protocol Availability
SPEED / SPEED+	Spacecraft 6D pose	RGB	Synthetic + lab	Pose labels, pose metrics	Public challenge splits
URSO / photoreal sets	Pose + tracking	RGB / stereo	Synthetic high-fidelity	Dense pose trajectories	Partial protocols
PhilEO Bench	EO segmentation/classification	Multi-spectral	Real satellite imagery	Pixel labels + scene tags	Benchmark suites available
Domain Adaptation RSI sets	EO transfer learning	RGB / multi-spectral	Mixed source/target	Class/pixel labels (subset)	Evaluation scripts in papers
Autonomy/GNC mission logs	Navigation / control	Vision + IMU	Real/sim mix	State trajectories, sparse events	Non-uniform; often custom
Onboard deployment testbeds	Runtime + efficiency	RGB / embedded streams	Hardware-in-loop	Throughput/latency logs	Limited standardization

Figure 3. Dataset and benchmark matrix for space-based perception themes, emphasizing realism, annotation depth, and reproducibility protocol maturity.

event streams [51]. These methods collectively ensure that AI models can handle the radiometric and geometric discrepancies inherent in the transition from simulation to the harsh reality of space [35, 51].

4. Novel Sensors and Multi-Modal Fusion

Beyond traditional RGB imaging, the emergence of event-based sensing and multi-modal fusion offers a resilient pathway to perception under extreme orbital illumination.

4.1. Event Cameras for Space Applications

Event-based cameras (EBCs) are a compelling alternative to traditional RGB sensors for spacecraft pose estimation, offering high temporal resolution ($\sim$1 μs), wide dynamic range (up to 140 dB), and low power consumption (0.5–1.5 W) [54]. Their high dynamic range is particularly advantageous in orbital environments, where extreme illumination and small solar exclusion angles often degrade conventional frame-based perception [54]. EBCs also exhibit inherent resilience to radiation-induced noise and, when coupled with Spiking Neural Networks (SNNs), can achieve up to 8$\times$ greater energy efficiency than standard architectures [11].

Recent research focuses on bridging the sim-to-real gap and developing efficient event representations. The SPADES dataset [54] provides a benchmark for these sensors, demonstrating that the asynchronous nature of EBCs can reduce the statistical discrepancy between synthetic and real-world distributions [30, 51]. For pose estimation, the 3-Channel (3C) representation outperforms alternatives like Event-to-Frame (E2F) or Time Surfaces (TS) by preserving temporal information through exponential decay across discrete sub-windows [54]. Architecturally, the Small Spiking End-to-End (S2E2) network [11] demonstrates fully event-based 6D pose estimation, achieving mean errors as low as 10.6 cm and 7.1° in cross-validation, making it suitable for resource-constrained on-orbit servicing and debris removal.

4.2. Multi-Modal and Non-Optical Sensing

To overcome the limitations of monocular RGB sensing, multi-modal fusion and non-optical sensors have gained significant traction. ChiNet [55] employs early-stage fusion by stacking thermal infrared (TIR) data with RGB inputs, utilizing a Recurrent Convolutional Neural Network (RCNN) to mitigate low-visibility artifacts and disambiguate thermal signatures from optical textures [13]. For non-cooperative targets, Inverse Synthetic Aperture Radar (ISAR) imaging provides a robust alternative; architectures like SU-Net [22] leverage dense residual blocks and multi-head self-attention to achieve state-of-the-art angular errors as low as 0.282° [62].

Active sensing via LiDAR and point cloud-based methods offer illumination-invariant 3D spatial data, critical for close-proximity operations [62]. Geometric attention mechanisms in Transformer-based architectures have been shown to outperform classical registration algorithms like ICP, achieving rotation errors of 0.492° and translation errors of 0.066 m. Stereo vision remains a vital low-SWaP solution for small satellite missions, with successful orbital validations on platforms such as SPHERES-VERTIGO and the Next Generation CubeSat Vision System [62]. For deep space exploration of minor bodies, hybrid techniques like stereophotoclinometry combine geometric and photometric data for high-fidelity terrain modeling and navigation [37].

5. Onboard AI and Edge Computing for Space

Deploying state-of-the-art AI in orbit necessitates a paradigm shift toward hardware-aware architectures that thrive within the severe power and radiation constraints of space-grade silicon.

Table 1. Key papers advancing Artificial Intelligence for Space from perception to autonomous operations.

Paper	Main Contribution
SPEED [30]	Foundational benchmark and dataset for monocular spacecraft pose estimation, bridging synthetic and real domains.
SDPENetv2 [73]	Achieves SOTA accuracy on SPEED using discrete pose weights and a learnable token head for efficient feature utilization.
UF-SPE [66]	Unified framework for cross-domain pose estimation integrating offline domain generalization with online unsupervised adaptation.
SPADES [54]	Introduces event-based sensing to space, providing a realistic dataset and 3-Channel representation to narrow the Sim2Real gap.
S2E2 [11]	Small Spiking End-to-End network demonstrating the feasibility of fully event-based, energy-efficient 6D pose estimation.
CogniSAT-6 [9]	Flight demonstration of dynamic targeting using onboard CNNs for real-time sensor redirection and anomaly detection.
PhilEO Bench [20]	Standardized evaluation framework for EO foundation models, revealing performance trade-offs between ViTs and U-Nets.
Bennunet [37]	Hybrid neural network architecture utilizing time-distributed CNNs and Keplerian dynamics for robust minor body navigation.
Edge SpAIce [14]	End-to-end pipeline for deploying optimized deep learning models on radiation-hardened satellite FPGAs using hls4ml.
CCT-Rel [4]	Characterizes radiation reliability of Transformers on edge TPUs, identifying Compact Convolutional Transformers as more robust.

5.1. Hardware Platforms for Edge Inference

Onboard processing requires specialized hardware to balance high throughput with stringent power constraints. The Coral Edge TPU has demonstrated significant efficiency for spacecraft pose estimation, achieving 7.7 FPS at 2.2W with quantized architectures, a 253% throughput increase and 25% power reduction over standard CPUs [39]. The Intel Myriad X VPU facilitates hybrid AI/CV pipelines by offloading classical vision tasks to SHAVE vector processors and neural inference to dedicated engines, maintaining approximately 5 FPS for 1MP images within a 2W envelope [34]. Neuromorphic processors like Intel's Loihi 2 offer transformative efficiency for event-driven tasks such as interference detection, outperforming conventional FPGAs by up to five orders of magnitude in energy-to-delay product [31]. NASA's SpaceCube platforms highlight these trade-offs: Zynq-PS SoCs provide maximum inference speed (up to 223× faster than softcore alternatives) at the cost of higher power, while RISC-V and MicroBlaze cores prioritize minimal power draw for less intensive missions [6]. A summary of key papers advancing AI for space is provided in Tab. 1.

5.2. Compression and Lightweight Architectures

Model compression and architectural optimization are essential for deploying deep neural networks on resource-constrained space platforms. Quantization to INT8 or Float16 significantly reduces model footprints; for instance, 8-bit quantization of MobileNetV3 achieves an 85.71% reduction in file size with a manageable 5% accuracy loss [2, 5]. To mitigate precision degradation, techniques like Quantization-Aware Training (QAT), Cross-Layer Equalization (CLE), and Fast Fine-Tuning (FFT) are frequently employed [23]. Lightweight architectures like SDPENet leverage depthwise separable convolutions and discrete Euler angle probability distributions to achieve up to 82.4% parameter reduction. These optimizations, combined with spatial-semantic interaction attention and multi-stage feature fusion, allow for significant improvements in pose estimation accuracy despite reduced model complexity [72].

5.3. Onboard Autonomy and Dynamic Targeting

Onboard autonomy enables satellites to react to transient events without ground intervention. The CogniSAT-6 mission demonstrates dynamic targeting by utilizing lookahead imagery and onboard CNNs to detect clouds or thermal anomalies, allowing for real-time sensor redirection within a 60–90 second window [9]. Such low-latency decision-making is supported by FPGA-accelerated frameworks; for instance, pruned U-Net models on Zynq UltraScale+ MPSoCs achieve cloud detection in under 30ms at 2.4W, while convolutional autoencoders provide 4× image compression with minimal latency [12]. To address sporadic connectivity and limited compute in LEO constellations, federated learning (FL) frameworks like FL-SEC and FEDGE have been proposed. These systems utilize'divide-and-conquer" strategies for model training and orbital aggregation protocols to reduce communication rounds, ensuring global model convergence despite intermittent visibility windows [17, 59].

6. Vision-Based Navigation for Space Missions

Vision-based navigation translates raw perceptual data into actionable trajectories, enabling precise landings and autonomous exploration across diverse celestial bodies.

6.1. Terrain Relative Navigation

Terrain Relative Navigation (TRN) is essential for absolute pose estimation in GPS-denied planetary environments. Systems like Image-Based Absolute Localization (IBAL) utilize satellite-derived templates and elevation maps as 3D landmarks [43]. Onboard imagery is matched to these templates via cross-correlation, followed by geometric verification using RANSAC and Perspective-n-Point (PnP) solvers [43]. At high altitudes, atmospheric distortion and horizon warping are mitigated through landmark angle filtering and incremental image corner projections [43]. Integration with gyroscopes enables forward propagation of pose estimates during rapid rotations or temporary landmark occlusion [43]. Flight demonstrations on suborbital platforms have achieved position accuracies ranging from sub-300 meters at 33 km to less than 55 meters at high velocities [43]. For

lunar missions, monocular TRN extracts crater centroids to match against 3D catalogs. Algorithms combining Efficient PnP (EPnP) with Iterative Closest Point (ICP) refinement have demonstrated sub-millimeter theoretical localization accuracy, though performance remains sensitive to lunar phase dynamics and orbital variations [27].

6.2. Minor Body and Deep Space Navigation

Autonomous exploration of small bodies requires multi-phase perception pipelines to handle unknown geometry and adverse lighting. During approach, centroid and light-curve algorithms estimate relative trajectories, while Shape-from-Silhouette (SfS) reconstructions determine initial visual hulls and pole orientations [47]. Proximity operations leverage optical flow for shape refinement and Stereo Photoclinometry (SPC) for dense 3D reconstruction of landing sites [47]. Frameworks employ covariance inflation and filter banks to manage system modeling errors and measurement transitions [47]. Recent hybrid neural architectures, such as Churinet and Bennunet, provide robust 6D pose estimation by combining CNN-based classification and regression [36, 37]. Bennunet further integrates time-distributed CNNs with Kalman filters based on Keplerian dynamics, significantly reducing tracking errors on real asteroid imagery [37]. These models are trained using GPU-accelerated physically based rendering (PBR) pipelines like SPyRender, which simulate complex illumination, albedo variations, and sensor noise to bridge the reality gap [36, 37].

6.3. Planetary Rover Navigation and Perception

Planetary rover autonomy relies on hierarchical GNC architectures that balance computational efficiency with navigation robustness. Modular graph-based SLAM frameworks, such as iSAM2, integrate local reference filters with 6D geometric features and ICP refinement to maintain global pose consistency [57]. These systems employ 3D probabilistic voxel-based maps (OctoMap) and perform loop closure via 6D map matching of geometric features [57]. Perception is enhanced by real-time semantic scene understanding using models like YOLOv5 for obstacle detection and stereo-depth estimation, with path planners employing A* algorithms on local semantic maps to generate unobstructed waypoints [56]. To address mobility challenges in soft terrain, vision-based methods employ machine learning (e.g., CNNs, Gaussian Processes) to predict wheel slip risk by classifying terrain types and estimating local slopes from 3D point cloud data [18]. Multi-modal fusion strategies combine inertial navigation (INS), wheel odometry, and VO, utilizing "zero-type updates" during periodic stops to constrain inertial drift and account for wheel slippage in unstructured, GPS-denied environments [29].

7. Robustness and Reliability

In the high-stakes domain of space, the transition from 'working' to 'trustworthy' AI requires rigorous frameworks for adversarial resilience and uncertainty quantification.

7.1. Adversarial Robustness and Security

Securing deep learning models against adversarial perturbations is critical for autonomous space operations. Wang et al. [64] introduced an adversarial attack detector for spacecraft pose estimation utilizing Long Short-Term Memory (LSTM) networks to monitor SHapley Additive exPlanations (SHAP) values. By extracting feature importance from the Global Average Pooling (GAP) layer, the detector identifies Fast Gradient Sign Method (FGSM) attacks with 99.21% accuracy in simulation and 96.29% in laboratory conditions [64]. Architectural resilience is also achieved through multi-sensor fusion and hybrid tracking systems combining KLT feature tracking with model-based approaches [62]. Input validation using binary classifiers ensures reliability by filtering unusable samples during eclipses or target absence [40]. Characterization of transformer models on edge TPUs reveals that Compact Convolution Transformers (CCT) offer higher radiation robustness than standard Vision Transformers (ViTs), which are sensitive to attention head counts [4].

7.2. Uncertainty Estimation and Calibration

Uncertainty quantification is fundamental to reliable autonomous navigation. Probabilistic models such as Mixture of Gaussians (MoG) have been employed for orientation soft classification, outperforming direct regression for non-cooperative spacecraft [52]. While Monte Carlo (MC) dropout remains a standard for estimating uncertainty in quaternion regression [52], discrete Euler angle probability distributions offer a lightweight alternative for attitude representation [72]. To bridge predictive reliability and hardware constraints, hardware-aware neural dropout search frameworks have been developed for FPGA deployment. These frameworks utilize evolutionary algorithms and Single Path One-Shot (SPOS) paradigms to optimize heterogeneous dropout configurations, achieving up to 33× higher energy efficiency than GPUs while maintaining well-calibrated uncertainty estimates [69]. Integrating these measures into adaptive filtering frameworks, such as the adaptive Unscented Kalman Filter (UKF), enhances robustness against dynamic maneuvers and variable lighting [60, 62].

7.3. Radiation Effects and Fault Tolerance

Radiation-induced faults are a primary challenge for edge computing in deep space. Experience from the OPS-SAT-1 mission highlighted the vulnerability of non-hardened eMMC NAND flash memory to file system corruption,

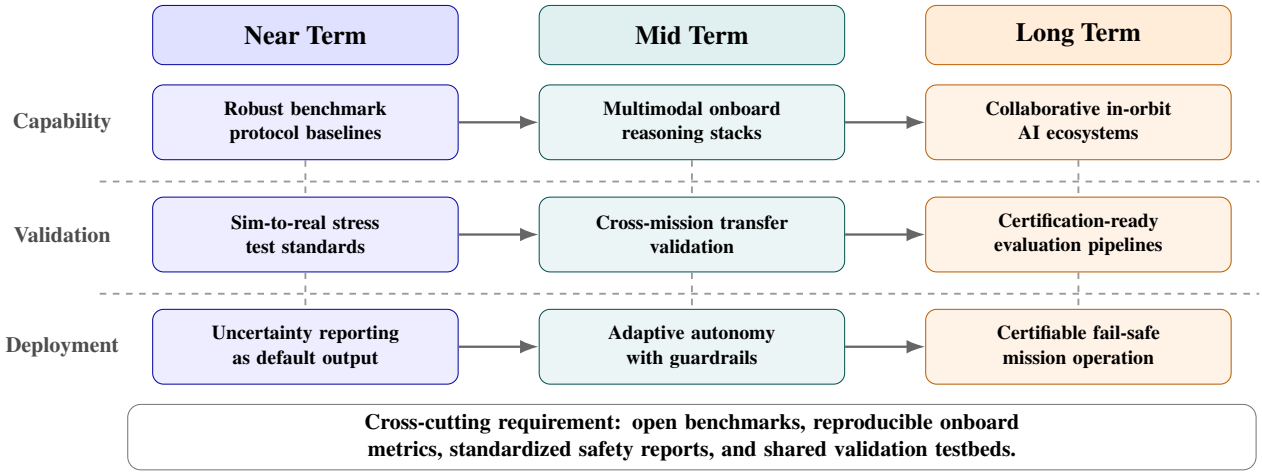

Figure 4. Research roadmap for autonomous space systems redesigned as a phased matrix, clarifying what matures first and what depends on prior validation.

necessitating volatile emergency modes (Ramfs) and "fully volatile" operational states to bypass permanent storage writes [26]. To mitigate these risks, space-qualified FPGAs and adaptive System-on-Chips (SoCs), such as the AMD Versal Edge, are preferred for deep learning acceleration due to their radiation tolerance and in-flight reconfigurability [12, 46]. These platforms enable hardware-level optimizations tailored to evolving mission requirements, such as cloud segmentation or image compression, while providing robust recovery through dynamic reconfiguration of processing pipelines [12, 26]. Adaptive SoCs allow missions to dynamically balance throughput and model complexity by adjusting memory architectures between local core memory and external DDR4, ensuring operational flexibility in resource-constrained environments [46].

8. Earth Observation and Remote Sensing

Onboard Earth observation is evolving from simple data collection to real-time intelligent synthesis, mitigating downlink bottlenecks through localized inference.

8.1. Onboard Processing for Earth Observation

Onboard processing is essential for Earth Observation (EO) to mitigate downlink bandwidth limitations and enable rapid response to time-critical events. The Edge SpAIce project exemplifies this trend, deploying deep convolutional neural networks (CNNs) on satellite Field Programmable Gate Arrays (FPGAs) using `hls4ml` and dataflow architectures [14]. To manage the resource constraints of space-grade hardware, optimization techniques including Knowledge Distillation (KD), Quantization-Aware Training (QAT) via the Brevitas library, and DSP packing are employed to maximize efficiency [14]. These methods have been validated for marine plastic litter detection using U-Net architectures on radiation-hardened FPGAs [14]. Wildfire detection systems utilize lightweight models such as VIB_SD and Refined UNet Lite, incorporating multispectral NIR and SWIR bands to maintain accuracy while minimizing computational overhead [5, 13]. Pruning and quantization remain critical for achieving real-time inference on edge accelerators like DPUs and FPGAs within strict Size, Weight, and Power (SWaP-C) envelopes [5, 12, 13].

8.2. Foundation Models for Earth Observation

The transition toward Self-Supervised Learning (SSL) and Foundation Models (FMs) aims to leverage massive unlabeled EO datasets. State-of-the-art models like Prithvi and SatMAE utilize Vision Transformer (ViT) and Masked Autoencoder (MAE) architectures to learn robust representations [20]. However, these models often struggle with fine-grained image-to-image tasks like segmentation due to latent space bottlenecks that discard low-level spatial information [20]. The PhilEO Bench framework provides standardized evaluation for Sentinel-2 FMs across tasks such as building density estimation and land cover classification, revealing that simpler U-Net architectures sometimes outperform ViT-based FMs in visual detail preservation [20]. Onboard deployment faces challenges including the exponential growth of multispectral data, a lack of diverse labeled datasets for rare events, and the discrepancy between ground-based training data and in-orbit sensor characteristics [12, 14, 17, 19, 23, 32, 61]. Robust cross-instrument generalization and hardware-aware design are required to ensure performance under strict SWaP-C constraints.

9. Open Challenges and Future Directions

The transition of deep learning (DL) models to flight missions is hindered by severe computational constraints, the sim-to-real domain gap, and rigorous verification requirements [4, 62]. Future research focuses on enhancing autonomy through Reinforcement Learning (RL) for fault-tolerant control and collaborative inference across satellite constellations [16, 67]. Adaptation of Foundation Models, such as LLMs and VLMs, enables robotic replanning and real-time fault diagnosis, though resource efficiency and hallucination rates remain significant hurdles [4, 45]. Digital Twins (DTs) and high-fidelity simulation frameworks like BSK-RL are vital for modeling operational constraints and clarifying trade-offs between simulation fidelity and in-orbit generalizability [45, 65]. Bridging these gaps requires hardware-software co-design, modular in-orbit updates, and specialized accelerators to ensure reliability under mission-critical tolerances [5, 15, 33].

To synthesize these challenges into an actionable trajectory, we propose a phased research roadmap for autonomous space perception (see Fig. 4). This roadmap delineates the progression from robust near-term benchmarking to the long-term goal of collaborative, certifiable in-orbit AI ecosystems, emphasizing the interdependence of capability, validation, and deployment.

10. Conclusion

The deployment of artificial intelligence in space is transitioning from experimental demonstrations to mission-critical operational capabilities. This review has synthesized the convergence of radiation-tolerant edge computing, event-driven perception, and domain-adaptive learning, highlighting the transformative impact of deep learning on spacecraft pose estimation, autonomous navigation, and onboard Earth observation. While significant progress has been made in developing lightweight architectures and multi-modal fusion strategies, such as the integration of event-based sensors and neuromorphic processors, substantial hurdles remain. The inherent sim-to-real domain gap necessitates continued innovation in high-fidelity simulation, self-supervised domain adaptation, and robust uncertainty quantification. Moreover, operating within the severe SWaP-C and radiation constraints of the space environment requires a paradigm shift toward hardware-software co-design and resilient edge computing. Addressing these open challenges will demand interdisciplinary collaboration across aerospace engineering, computer vision, and hardware architecture. Ultimately, the maturation of verifiable, fault-tolerant AI systems will pave the way for a new era of autonomous space exploration, enabling responsive satellite constellations and intelligent deep-space missions.

Acknowledgements

The intellectual organization, critical synthesis, and structured narrative of this review were entirely conceived and developed by the authors. To assist in retrieving and extracting factual information from the surveyed papers, the authors utilized PaperQA [41], a Retrieval-Augmented Generation (RAG) tool for scientific literature. All retrieved information was carefully verified and contextualized within the authors' original review framework.

References

[1] Louis Aberdeen, M. Hansen, Melvyn L. Smith, and Lyndon N. Smith. Deep learning-based image recovery and pose estimation for resident space objects. In *arXiv.org*, 2025. 1

[2] W. Akeboshi, R. Billones, Jerahmeel K. Coching, Adrian Jenssen Lim Pe, Seth Gabriel Dy Yeung, R. J. Tan Ai, E. Sybingco, Elmer P. Dadios, and Mark Angelo C. Purio. Evaluation of quantized cnn architectures for land use classification for onboard cube satellite computing. In *International Conference on Optoelectronics and Microelectronics*, 2024. 5

[3] N. Ammour, Laila Bashmal, Y. Bazi, M. M. Rahhal, and M. Zuair. Asymmetric adaptation of deep features for cross-domain classification in remote sensing imagery. In *IEEE Geoscience and Remote Sensing Letters*, 2018. 1

[4] Pablo R. Bodmann, Paolo Rech, and Matteo Saveriano. Evaluating the reliability of vision transformers for space robotics applications. In *2024 International Conference on Space Robotics (iSpaRo)*, 2024. 5, 6, 8

[5] Anfal Yousra Boucetta, Meriem Baziz, Leila Hamdad, and Imed Allal. Optimizing for edge-ai based satellite image processing: A survey of techniques. In *2024 IEEE Mediterranean and Middle-East Geoscience and Remote Sensing Symposium (M2GARSS)*, 2024. 5, 7, 8

[6] T. Chase, Justin Goodwill, Karthik Dantu, and Christopher Wilson. Profiling vision-based deep learning architectures on nasa spacecube platforms. In *IEEE Aerospace Conference*, 2024. 1, 5

[7] Bo Chen, Jiewei Cao, Álvaro Parra, and Tat-Jun Chin. Satellite pose estimation with deep landmark regression and nonlinear pose refinement. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 2019. 2

[8] Zilong Chen, Shijie Wang, Haichao Gui, and Rui Zhong. Deep-learning-based pose uncertainty estimation for noncooperative spacecraft with unit dual quaternion distributions. In *Acta Astronautica*, 2026. 1

[9] Steve A. Chien, Itai Zilberstein, Alberto Candela, David Rijlaarsdam, Tom Hendrix, Aubrey Dunne, Oriol Aragon Casaled, and Juan Puig Miquel. Flight of dynamic targeting on the cognisat-6 spacecraft. In *arXiv.org*, 2025. 5

[10] Kiruki Cosmas and Asami Kenichi. Utilization of fpga for onboard inference of landmark localization in cnn-based spacecraft pose estimation. In *Aerospace*, 2020. 1

[11] Jonathan Courtois, Benoît Miramond, and Alain Pegatoquet. Spiking monocular event based 6d pose estimation for space application. In *arXiv.org*, 2025. 4, 5

[12] Angela Cratere, Marcello Asciolla, and Francesco Dell'Olio. Towards the integration of fpga-based deep learning edge computing on smallsats for low-latency autonomous decision-making. In *Sensors Applications Symposium*, 2025. 5, 7

[13] Khushal Das, Jiju Poovvancheri, Sergio Flesca, Claudia Roberta Calidonna, and Dong Chen. Emerging trends in wildfire detection through the lens of computer vision and wildfire emission quantification: A comprehensive survey. In *IEEE Access*, 2026. 4, 7

[14] Noemi D'Abbondanza, Stylianos Tzelepis, Nicolò Ghielmetti, Ioannis Kakogeorgiou, Vanya Buchova, Konstantinos Karantzalos, K. Kikaki, Nicolas-Marcel Lemoine, M. Pierini, S. Summers, Simon Vellas, F. D. Vieilleville, and Boyan-Nikola Zafirov. Edge spaice: Deep learning deployment pipeline for onboard data reduction on satellite fpgas. In *IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum*, 2025. 5, 7

[15] Jonah Ekelund, Ricardo Vinuesa, Yuri V. Khotyaintsev, Pierre Henri, G. Delzanno, and Stefano Markidis. Ai in space for scientific missions: Strategies for minimizing neural-network model upload. In *IEEE International Conference on e-Science*, 2024. 8

[16] Ghaith El-Dalahmeh, M. Sattari, Q. Vo, and Ryszard Kowalczyk. Intelligent control of spacecraft reaction wheel attitude using deep reinforcement learning. In *arXiv.org*, 2025. 8

[17] Mohamed Elmahallawy and Tie Luo. Stitching satellites to the edge: Pervasive and efficient federated leo satellite learning. In *Annual IEEE International Conference on Pervasive Computing and Communications*, 2024. 5, 7

[18] Masafumi Endo, S. Endo, K. Nagaoka, and Kazuya Yoshida. Terrain-dependent slip risk prediction for planetary exploration rovers. In *Robotica (Cambridge. Print)*, 2021. 6

[19] Li Fang, Tianyu Li, Yanghong Lin, Shudong Zhou, and Wei Yao. Ai-enabled satellite edge computing: A single-pixel feature based shallow classification model for hyperspectral imaging. In *arXiv.org*, 2026. 7

[20] Casper Fibaek, Luke Camilleri, Andreas Luyts, Nikolaos Dionelis, and B. L. Saux. Phileo bench: Evaluating geospatial foundation models. In *IEEE International Geoscience and Remote Sensing Symposium*, 2024. 1, 5, 7

[21] Fabien Gallet, Christophe Marabotto, and Thomas Chambon. Exploring ai-based satellite pose estimation: from novel synthetic dataset to in-depth performance evaluation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 3

[22] Huiyu Gao, Zhihui Li, Ning Wang, Jingfan Yang, and Depeng Dang. Su-net: pose estimation network for non-cooperative spacecraft on-orbit. In *Scientific Reports*, 2023. 3, 4

[23] L. M. Garcés-Socarrás, Raudel Cuiman, Flor Ortíz, J. A. Vásquez-Peralvo, J. L. González-Rios, Mouhamad Chehaitly, Arkadii Kazanskii, Sahar Malmir, Amirhossein Nik, Jan Thoemel, Sumit Kumar, Marcele Kuhfuss, Swetha Varadajulu, Eva Lagunas, JUAN CARLOS MERLANO DUNCAN, Jorge Querol, and S. Chatzinotas. Onboard machine learning for satellite edge computing: The spaice project use case. In *2025 European Data Handling & Data Processing Conference (EDHPC)*, 2025. 5, 7

[24] A. García, M. A. Musallam, Vincent Gaudillière, Enjie Ghorbel, Kassem Al Ismaeil, M. Perez, and D. Aouada. Lspnet: A 2d localization-oriented spacecraft pose estimation neural network. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021. 1, 2

[25] Mohammad Tanvir Mahmud Habib and Shahnewaz Siddique. Satellite pose estimation using attention learning without gpu acceleration. In *2025 International Conference on Emerging Techniques in Computational Intelligence (ICETCI)*, 2025. 3

[26] Maximilian Henkel, Vladimir Zelenevskiy, Georges Labrèche, Rodrigo Laurinovics, David Evans, Dominik Marszk, and Omiros Papadatos Vasilakis. Mitigating and recovering from radiation induced faults in non-hardened spacecraft flash memory. In *IEEE Aerospace Conference*, 2024. 7

[27] P. Hu, Bo Li, and D. Zou. Lunar satellite position estimation based on monocular vision localization. In *2025 International Conference on Future Communications and Networks (FCN)*, 2025. 6

[28] Y. Huo, Zhi Li, and Feng Zhang. Fast and accurate spacecraft pose estimation from single shot space imagery using box reliability and keypoints existence judgments. In *IEEE Access*, 2020. 2

[29] Cagri Kilic, Jason N. Gross, Nicholas Ohi, Ryan M. Watson, Jared Strader, Tommy Swiger, Scott Harper, and Yu Gu. Improved planetary rover inertial navigation and wheel odometry performance through periodic use of zero-type constraints. In *IEEE/RJS International Conference on Intelligent RObots and Systems*, 2019. 6

[30] Mate Kisantal, Sumant Sharma, T. Park, D. Izzo, Marcus Märtens, and S. D'Amico. Satellite pose estimation challenge: Dataset, competition design, and results. *IEEE Transactions on Aerospace and Electronic Systems*, 2019. 1, 2, 3, 4, 5

[31] Eva Lagunas, Flor G. Ortiz-Gomez, G. Eappen, Saed Daoud, W. A. Martins, J. Querol, S. Chatzinotas, N. Skatchkovsky, Bipin Rajendran, and Osvaldo Simeone. Performance evaluation of neuromorphic hardware for onboard satellite communication applications. In *IEEE wireless communications*, 2024. 5

[32] Jérémy Lebreton, Ingo Ahrns, R. Brochard, Christoph Haskamp, M. L. Goff, Nicolas Menga, Nicolas Ollagnier, Ralf Regele, Francesco Capolupo, and M. Casasco. Training datasets generation for machine learning: Application to vision based navigation. In *arXiv.org*, 2024. 3, 7

[33] G. Lentaris, I. Stratakos, I. Stamoulias, D. Soudris, Manolis I. A. Lourakis, and Xenophon Zabulis. High-performance vision-based navigation on soc fpga for spacecraft proximity operations. *IEEE transactions on circuits and systems for video technology (Print)*, 2020. 1, 8

[34] Vasileios Leon, Panagiotis Minaidis, G. Lentaris, and Dimitrios Soudris. Accelerating ai and computer vision for satellite pose estimation on the intel myriad x embedded soc. In *Microprocessors and microsystems*, 2023. 5

[35] Kun Liu and Yongjun Yu. Revisiting the domain gap issue in non-cooperative spacecraft pose tracking. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 2, 3, 4

[36] A. E. Lopez, Pablo Ghiglino, and M. Sanjurjo-Rivo. Churinet—applying deep learning for minor bodies optical navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 2023. 3, 6

[37] A. E. Lopez, Pablo Ghiglino, and M. Sanjurjo-Rivo. Bennunet—an update on applying deep learning for minor bodies optical navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 2025. 2, 3, 4, 5, 6

[38] Alessandro Lotti, D. Modenini, P. Tortora, M. Saponara, and M. Perino. Deep learning for real time satellite pose estimation on low power edge tpu. In *arXiv.org*, 2022. 1

[39] Alessandro Lotti, D. Modenini, P. Tortora, M. Saponara, and M. Perino. Deep learning for real-time satellite pose estimation on tensor processing units. *Journal of Spacecraft and Rockets*, 2023. 1, 5

[40] Lucrezia Lovaglio and Fabrizio Stesina. Deep learning-optimized monocular navigation for autonomous rendezvous and proximity maneuvers in small satellite missions. In *IEEE International Workshop on Metrology for AeroSpace*, 2025. 1, 3, 6

[41] Jakub Lála, Odhran O'Donoghue, Aleksandar Shtedritski, Sam Cox, Samuel G. Rodriques, and Andrew D. White. Paperqa: Retrieval-augmented generative agent for scientific research. *arXiv preprint arXiv:2312.07559*, 2023. 8

[42] Xianping Ma, Xiaokang Zhang, Xingchen Ding, M. Pun, and Siwei Ma. Decomposition-based unsupervised domain adaptation for remote sensing image semantic segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 2024. 1

[43] Dominic Maggio, C. Mario, Brett J. Streetman, T. Steiner, and L. Carlone. Vision-based terrain relative navigation on high-altitude balloon and sub-orbital rocket. In *AIAA SCITECH 2023 Forum*, 2023. 5

[44] Trupti Mahendrakar, Ryan T. White, M. Wilde, and Madhur Tiwari. Spaceyolo: A human-inspired model for real-time, on-board spacecraft feature detection. In *IEEE Aerospace Conference*, 2023. 1

[45] Nikos Mavrakis, Lai-Chong Law, Hubert P. H. Shum, Leonie Bensch, Tommy Nilsson, Martin Nisser, Pat Pataranutaporn, Albrecht Schmidt, and Valentina Sumini. Integrating human-in-the-loop ai to tackle space communication delay challenges. In *SpaceCHI*, 2025. 1, 8

[46] Andrew McCormick, Adewale Adetomi, and Ken O'Neill. Deploying ai-based spacecraft telemetry anomaly detection on an adaptive system-on-chip solution. In *2025 European Data Handling & Data Processing Conference (EDHPC)*, 2025. 7

[47] I. Nesnas, Benjamin J. Hockman, Saptarshi Bandopadhyay, B. Morrell, D. Lubey, Jacopo Villa, D. Bayard, Alan Osmundson, Benjamin Jarvis, Michele Bersani, and S. Bhaskaran. Autonomous exploration of small bodies toward greater autonomy for deep space missions. In *Frontiers in Robotics and AI*, 2021. 1, 3, 6

[48] Mi Ning, Shijie Zhang, and Shiqiang Wang. A non-cooperative satellite feature point selection method for vision-based navigation system. In *Italian National Conference on Sensors*, 2018. 1

[49] Leo Pauly, Wassim Rharbaoui, C. Shneider, A. Rathinam, Vincent Gaudillière, and D. Aouada. A survey on deep learning-based monocular spacecraft pose estimation: Current state, limitations and prospects. In *Acta Astronautica*, 2023. 1

[50] M. Piazza, Michele Maestrini, and P. Lizia. Deep learning-based monocular relative pose estimation of uncooperative spacecraft, 2021. 2, 3

[51] Mirco Planamente, Chiara Plizzari, Marco Cannici, Marco Ciccone, Francesco Strada, A. Bottino, M. Matteucci, and Barbara Caputo. Da4event: Towards bridging the sim-to-real gap for event cameras using domain adaptation. In *IEEE Robotics and Automation Letters*, 2021. 4

[52] Pedro F. Proença and Yang Gao. Deep learning for spacecraft pose estimation from photorealistic rendering. In *IEEE International Conference on Robotics and Automation*, 2019. 1, 2, 3, 6

[53] Roman Prokazov, Alessandro Lotti, D. Modenini, M. R. Makhdoomi, Carol Martínez, M. Olivares-Mendez, G. Kremer, and P. Tortora. Deep learning-based spacecraft pose estimation with a fail-safe fallback mechanism. In *23rd IAA Symposium on Space Debris*, 2025. 1

[54] A. Rathinam, Haytam Qadadri, and D. Aouada. Spades: A realistic spacecraft pose estimation dataset using event sensing. In *IEEE International Conference on Robotics and Automation*, 2023. 3, 4, 5

[55] Duarte Rondao, N. Aouf, and Mark A. Richardson. Chinet: Deep recurrent convolutional learning for multimodal spacecraft pose estimation. *IEEE Transactions on Aerospace and Electronic Systems*, 2021. 4

[56] Ragav Sachdeva, Ravi Hammond, J. Bockman, Alec Arthur, Brandon Smart, Dustin Craggs, Anh-Dzung Doan, T. Rowntree, Elijah Schutz, Adrian Orenstein, Andy Yu, Tat-Jun Chin, and I. Reid. Autonomy and perception for space mining. In *IEEE International Conference on Robotics and Automation*, 2021. 6

[57] M. J. Schuster, S. Brunner, Kristin Bussmann, Stefan Büttner, Andreas Dömel, Matthias Hellerer, H. Lehner, Peter Lehner, O. Porges, Josef Reill, Sebastian Riedel, Mallikarjuna Vayugundla, B. Vodermayer, T. Bodenmüller, Christoph Brand, W. Friedl, I. Grixa, H. Hirschmüller, M. Kassecker, Zoltán-Csaba Márton, Christian Nissler, Felix Ruess, M. Suppa, and A. Wedler. Towards autonomous planetary exploration, 2019. 1, 6

[58] Dhruv Shah and S. Levine. Viking: Vision-based kilometer-scale navigation with geographic hints. In *Robotics: Science and Systems*, 2022. 1

[59] Rashmi Siddalingappa, Deepa S, Priya Stella Mary I, K. P, and Lakshmi B A. Fedge: Federated learning at the edge

on space platforms using deep neural network architectures. *International journal of information technology*, 2025. 5

[60] J. Sosa, Dan Pineau, A. Rathinam, Abd El Rahman Shabayek, and D. Aouada. Motion aware vit-based framework for monocular 6-dof spacecraft pose estimation. In *arXiv.org*, 2025. 2, 3, 6

[61] Bernát-Kristóf Vescan-Bakcsy and Atena Stefania Szabo. The great coral reef: Modular ai/ml task acceleration for onboard data processing. In *2025 European Data Handling & Data Processing Conference (EDHPC)*, 2025. 7

[62] Kenny Vishal, Vishnu G. Nair, and Jeane Marina D'Souza. Review of spacecraft pose-estimation techniques for close-proximity operations. In *IEEE Access*, 2025. 1, 2, 3, 4, 6, 8

[63] Jingrong Wang, Bo Xu, Rong Jin, Shoujian Zhang, Kefu Gao, and Jingnan Liu. Sky-gvio: an enhanced gnss/ins/vision navigation with fcn-based sky-segmentation in urban canyon. In *Remote Sensing*, 2024. 1

[64] Ziwei Wang, Nabil Aouf, Jose Pizarro, and Christophe Honvault. Robust adversarial attacks detection for deep learning based relative pose estimation for space rendezvous. In *Advances in Space Research*, 2023. 6

[65] Cannon Whitney and J. Melville. Toward trusted onboard ai: Advancing small satellite operations using reinforcement learning. In *arXiv.org*, 2025. 1, 8

[66] Yingjian Yu, Zhang Li, and Qifeng Yu. A unified framework for cross-domain space drone pose estimation integrating offline domain generalization with online domain adaptation. In *Drones*, 2025. 5

[67] Honggang Yuan, Zixin Wang, Yuning Jiang, Xin Liu, Yuanming Shi, and Ting Wang. Sai: Latency-aware satellite edge lam inference with looped transformer. In *ICC 2025 - IEEE International Conference on Communications*, 2025. 8

[68] Yunfeng Zhang, Linjing You, Luyu Yang, Zhiwei Zhang, Xiangli Nie, and Bo Zhang. Monocular 6-dof pose estimation of spacecrafts utilizing self-iterative optimization and motion consistency. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 2, 3

[69] Zehuan Zhang, Hongxiang Fan, H. Chen, Lukasz Dudziak, and Wayne Luk. Hardware-aware neural dropout search for reliable uncertainty prediction on fpga. In *Design Automation Conference*, 2024. 6

[70] Juepeng Zheng, Yi Zhao, Wenzhao Wu, Mengxuan Chen, Weijia Li, and H. Fu. Partial domain adaptation for scene classification from remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 2023. 1

[71] Hang Zhou, Lu Yao, Haoping She, and Weiyong Si. Sdpenet: A lightweight spacecraft pose estimation network with discrete euler angle probability distribution. In *IEEE Robotics and Automation Letters*, 2025. 1

[72] Hang Zhou, Lu Yao, Haoping She, and Weiyong Si. Sdpenet: A lightweight spacecraft pose estimation network with discrete euler angle probability distribution. In *IEEE Robotics and Automation Letters*, 2025. 2, 3, 5, 6

[73] Hang Zhou, Lu Yao, Haoping She, and Weiyong Si. Sdpenetv2: Spacecraft pose estimation network with learnable token head based on discrete pose weights. In *IEEE Robotics and Automation Letters*, 2025. 2, 5

[74] Jingru Zhu, Ya Guo, Geng Sun, Libo Yang, M. Deng, and Jie Chen. Unsupervised domain adaptation semantic segmentation of high-resolution remote sensing imagery with invariant domain-level prototype memory. *IEEE Transactions on Geoscience and Remote Sensing*, 2022. 1