Journal Homepage: www.ijarpr.com ISSN: 3049-0103 (Online)

International Journal of Advance Research Publication and Reviews

Vol 03, Issue 9, pp 222-232, September, 2026

Emerging Satellite Technologies for Next-Generation Earth Observation and Navigation *A Review of Sensor Fusion, Artificial Intelligence, and Architectural Innovations* A Synthesis of 23 Studies on Sensor Fusion, Navigation, and Positioning (2024–2026)

[1]Wash Patrick Madugu, [2]Otu Simon Achu, [3]Mustafa Ibrahim Ahmed, [4]Amanda-Roy C. Oguejiofo, [5]Bello Muazu Maccido, [6]Safwan Sani Toro, [7]Awobotu Olaide Marcus

[1-7]Department of Mission Planning and Satellite Data Management, National Space Research and Development Agency (NASRDA).

ABSTRACT :

Context: Global Navigation Satellite Systems (GNSS) remain the backbone of modern positioning, yet their fragility in urban canyons, forests, tunnels, and jammed environments has catalyzed a broad shift toward multi-sensor fusion architectures. Scope: This paper synthesizes 23 peer-reviewed studies published between 2024 and 2026, spanning precision agriculture, wildfire and flood mapping, GNSS-denied navigation, autonomous robotics, on-orbit satellite computing, and integrated satellite communication-navigation (ICAN) systems. Method: Through thematic synthesis and comparative analysis, we identify three converging threads: (i) the retirement of single-sensor systems in favor of redundant satellite-aerial-ground pipelines; (ii) the gradual displacement of classical Kalman filtering by learned fusion operators including graph neural networks, attention mechanisms, and deep reinforcement learning; and (iii) the emergence of LEO mega-constellations as dual-purpose communication-navigation infrastructure. Insight: We propose a six-layer reference architecture (sensing, edge and on-orbit pre-processing, adaptive fusion, AI-driven intelligence, communication and networking, and application-level decision-making), tied together by a continuous feedback loop rather than a one-way pipeline. The review closes with a gap-frequency analysis, a technology-readiness assessment, and an honest appraisal of persistent unsolved problems:

cold-start sensitivity, on-board computational limits, non-line-of-sight signal degradation, and the absence of shared cross-domain benchmarks.

Keywords: *sensor fusion; GNSS-denied navigation; LEO constellations; integrated communication and navigation; Earth observation; deep learning; Kalman filter; on-orbit computing*

1. INTRODUCTION

Modern life depends on knowing where things are. A tractor needs centimeter-level accuracy to avoid re-spraying the same crop row. A delivery drone must avoid power lines it cannot see. A self-driving car must localize itself inside a parking garage where satellites cannot reach. For decades, the default answer to "where am I?" has been GNSS: GPS, Galileo, BeiDou, and GLONASS. It is remarkable engineering, but it has a well-known Achilles' heel: it requires a clear line of sight to satellites tens of thousands of kilometers away. That makes it fragile exactly where demand is highest, in dense cities, forests, tunnels, indoor spaces, and any environment where jamming or spoofing is present.

The response has not been a single breakthrough technology. It has been a broad, improvisational combination of approaches. On one front, researchers fuse GNSS with inertial sensors, cameras, and LiDAR so that when the satellite signal drops, the system coasts on other information rather than losing track entirely. On another front, researchers ask whether LEO broadband constellations, such as Starlink, OneWeb, and Kuiper, can broadcast navigationgrade signals as a secondary service, since their proximity to Earth yields signals 20–30 dB stronger than traditional GNSS. A third group looks past positioning altogether, fusing satellite and drone imagery to answer questions about crop health, flood extent, wildfire risk, and urban growth that no ground survey could match in speed or scale.

These efforts rarely cite each other. They sit in different sub-fields with different vocabularies. Read together, however, they tell a coherent story about where satellite-enabled sensing is headed. This paper reviews 23 studies published in 2024–2026 and asks two questions. First, what do these studies have in common at the level of actual engineering choices? Second, if those commonalities were distilled into a single reference architecture, what would it look like, and how would data move through it from orbit to decision?

The remainder of the paper is organized as follows. Section 2 describes the review protocol. Section 3 thematically analyzes the literature across five domains. Section 4 presents a comparative summary table and quantitative meta-figures. Section 5 proposes the six-layer architectural framework. Section 6 maps research gaps through a frequency matrix and technology-readiness assessment. Section 7 acknowledges the limitations of this review. Section 8 concludes.

2. REVIEW METHODOLOGY

2.1 Search Protocol and Inclusion Criteria

This review follows a systematic literature review protocol adapted from PRISMA guidelines. We queried IEEE Xplore, Scopus, Web of Science, and arXiv for peer-reviewed articles and major conference proceedings published between January 2024 and August 2026. The search strings combined three concept blocks:

- Block A (sensing platforms): ("satellite" OR "UAV" OR "drone" OR "LEO constellation" OR "GNSS")
- Block B (fusion/intelligence): ("sensor fusion" OR "data fusion" OR "deep learning" OR "Kalman filter" OR "attention mechanism" OR "graph neural network")
- Block C (application): ("navigation" OR "positioning" OR "Earth observation" OR "remote sensing" OR "autonomous")

Inclusion criteria were: (i) peer-reviewed journal article or major conference proceedings (IROS, CVPR, ICASSP); (ii) primary focus on multi-sensor fusion, satellite-based sensing, or navigation and positioning; (iii) quantitative results reported; (iv) English language. Exclusion criteria were: (i) pure hardware design without fusion methodology; (ii) non-sensor aerial platforms only (e.g., manned aircraft without satellite integration); (iii) editorial or commentary pieces without original experiments.

2.2 Screening and Selection

The initial query returned 1,847 records. After removing duplicates (n = 312), title and abstract screening excluded 1,398 records that did not meet the inclusion criteria. Full-text review of the remaining 137 articles led to the final set of 23 studies that form the core of this synthesis. Figure 2a summarizes coverage by application domain; Figure 2b shows the distribution of fusion methods.

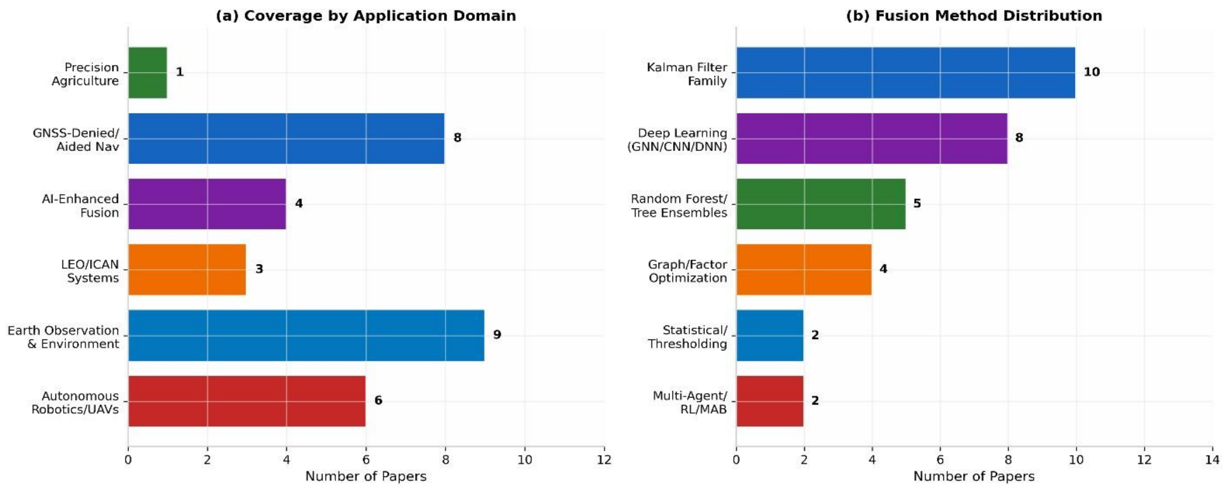

Figure 2. (a) Coverage of the 23 reviewed studies by application domain. (b) Distribution of fusion methods used across the reviewed literature.

Earth observation and environmental monitoring is the largest single domain, accounting for 9 of the 23 studies, followed by GNSS-denied and GNSSaided navigation (8 studies) and autonomous robotics or UAV navigation (6 studies). AI-enhanced fusion and LEO or ICAN systems are smaller but growing clusters, at 4 and 3 studies respectively, while precision agriculture, though the subject of only 1 dedicated study here, supplies one of the review's clearest fusion case studies (Avioz et al., 2025). On the methods side, the classical Kalman filter family remains the single most common fusion approach, appearing in 10 of the 23 studies, ahead of deep-learning methods such as GNNs, CNNs, and DNNs (8 studies) and tree-ensemble methods like Random Forest and XGBoost (5 studies). Graph and factor-optimization approaches appear in 4 studies, and statistical thresholding and multi-agent or reinforcement-learning methods each appear in 2. Because several studies combine more than one fusion family, as Section 3 discusses, the two panels of Figure 2 do not sum neatly to 23, and a paper counted under Kalman filtering may also appear under deep learning if it uses both.

2.3 Data Extraction

For each selected study, we extracted: (i) sensors and platforms used; (ii) fusion or learning methodology; (iii) headline quantitative results; (iv) authoridentified gaps; and (v) broader claims about the field. That extraction exercise produced the comparative table in Section 4 and surfaced the recurring architectural pattern formalized in Section 5.

3. THEMATIC REVIEW OF THE LITERATURE

Reading all 23 studies together, five natural groupings emerge. Several papers, such as Sharma (2026) and Ekolama (2026), sit comfortably in two categories at once, reflecting the interdisciplinary nature of the field.

3.1 GNSS-Denied and GNSS-Aided Positioning

The most direct response to GNSS vulnerabilities is to stop relying on it alone. Susanti et al. (2026) and Wang and Ahmad (2025) both conducted largescale, bibliometric-informed surveys and reached strikingly consistent conclusions: combining GPS with IMU, LiDAR, and vision can shrink localization error from roughly 79 meters down to 3.7 meters. Both surveys found the classical Kalman filter family (EKF, UKF, CKF) still doing most of the heavy lifting, even as AI-based approaches earn their place alongside them. Susanti et al. (2026), drawing on a bibliometric review of work from 2018 to 2024, further report that a cubature Kalman filter (CKF) variant improves on plain GPS accuracy by 34%, underscoring how much of that 79 m to 3.7 m gain still comes from refinements within the Kalman family rather than from replacing it outright.

Jarraya et al. (2025) propose a cleaner taxonomy: they split methods into "absolute" techniques, which match sensor observations to external references such as terrain maps (TERCOM, SITAN, DSMAC), and "relative" techniques, which track movement from a known starting point using visual-inertial odometry or SLAM. Their conclusion, echoed across the other GNSS-denied studies, is that no single sensor suffices; the honest engineering answer is always hybrid.

Sharma (2026) shows that matching a drone's downward-facing camera to satellite imagery can hold position error under three meters even at 60–150 m altitude, a 3.5× improvement over VIO alone. Alghamdi et al. (2025) look further afield at jamming-resistant alternatives, including celestial navigation, quantum sensing, and terrestrial radio beacons, that could serve as a last line of defense when every other signal has failed; in their review, a quantumnavigation system reportedly sustained a four-hour flight without any GPS input at all. Feng et al. (2025) tackle the safety-critical case of autonomous driving, building a noise-adaptive ESKF that improves uncertainty quantification by 37.7% over elevation-based methods.

3.2 AI-Enhanced and Learned Fusion Techniques

A second group of papers is less interested in which sensors to combine and more interested in changing how the combining itself is done. For sixty years, the Kalman filter and its descendants have been the default mathematical machinery for fusion. They are well understood, computationally cheap, and provably optimal under Gaussian assumptions. Real sensor noise rarely obeys those assumptions, and this is the gap that newer AI-based approaches fill.

Ekolama (2026) combines multi-constellation GNSS, LEO signals, and situational context through a graph neural network with an attention mechanism. The resulting system holds accuracy to a median of 6.2 centimeters in dense urban settings, beating a conventional EKF baseline by 42–57%. Yuan, Yu, and Zong (2025) take a more radical step, training an artificial neural network to learn fusion rules for launch-vehicle navigation directly from flight data. The resulting system outperforms a UKF on both position (9.1 m vs. 67.4 m RMSE) and velocity (0.015 m/s vs. 0.17 m/s RMSE). Guo et al. (2025) push learned fusion into 3D scene understanding: their SGFormer network uses satellite imagery to correct the blind spots and occlusions that a ground-level camera alone cannot see, achieving 45.01 IoU on SemanticKITTI.

A growing insistence on explainability runs through this group. Tools like SHAP appear repeatedly: Dong and Yuan (2026) on soil moisture, Al Saim and Aly (2025) on tree-species mapping, and Durlevic, Ilic, and Valjarevic (2025) on wildfire risk, so that engineers can see which sensor or feature drives a given prediction rather than trusting a black box with a safety-critical decision. Dong and Yuan's (2026) SHAP analysis of their soil-moisture fusion model, for instance, attributes roughly 66% of predictive weight to static terrain parameters and 20% to the incoming satellite soil-moisture signal itself, while an XGBoost-guided feature-elimination step trims 57% of candidate inputs without hurting accuracy. Al Saim and Aly (2025) similarly find that elevation, ahead of vegetation indices, is the single most influential variable in their tree-species classifier.

3.3 LEO Constellations and Integrated Communication-Navigation Systems

A third thread has almost nothing to do with sensor fusion in the traditional sense and everything to do with infrastructure economics. Thousands of LEO satellites have already been launched for broadband internet. A growing body of research asks whether that same hardware could double as a navigation system almost for free.

Wang et al. (2025b) lay out the case clearly: because LEO satellites orbit much closer to Earth than traditional GNSS satellites, their signals arrive 20–30 dB stronger. That stronger signal cuts precise-point-positioning convergence from roughly thirty minutes down to under three. Wang et al. (2025a) tackle the harder engineering problem underneath: if satellites are to do more processing themselves rather than shipping raw data to the ground, they need computing hardware that survives in orbit. Their review of hyper-heterogeneous on-board architectures shows efficiency levels now exceeding one TOPS per watt. Hashima et al. (2025) complete the picture from the communications side, surveying how AI, and multi-armed bandit algorithms in particular, can allocate scarce bandwidth between UAVs and satellites in a 6G network, reporting energy-efficiency gains of 69 to 118% and a fairness index roughly 87 times better than fixed-allocation benchmarks; across their broader survey, deep reinforcement learning methods account for about a third of all AI techniques proposed for this problem.

Read together, these three papers describe an industry converging on satellites that talk, compute, and navigate all at once.

3.4 Earth Observation for Agriculture and Environmental Monitoring

The largest single cluster of studies has nothing to do with knowing where a vehicle is and everything to do with knowing what is happening on the ground.

Avioz et al. (2025) offer the most tightly scoped example: by fusing drone multispectral imagery, Sentinel-2 passes, and canopy structure from structurefrom-motion photogrammetry, they predict nitrogen content in citrus leaves well enough ($R^2 = 0.80$) to show a usable link between a tree's nitrogen levels and its eventual yield. The same fusion logic scales to larger problems. Durlevic, Ilic, and Valjarevic (2025) combine four satellite data sources with deep learning to map wildfire risk across Serbia. Fawakherji and Hashemi-Beni (2025) pair drone optical imagery with Sentinel-1 radar to map flood extent with 94.12% accuracy; radar is essential because it sees through the cloud cover that floods usually bring.

A similar "optical plus radar plus machine learning" pattern recurs across surface-water monitoring (Declaro and Kanae, 2024), soil-moisture mapping across China (Dong and Yuan, 2026), urban expansion tracking in southern Italy (Vitale and Lamonaca, 2025), tree-species classification in Arkansas (Al Saim and Aly, 2025), and ocean current reconstruction in the Mediterranean (Kugusheva et al., 2024). Wang et al. (2025c) apply multi-source satellite fusion (ZY-3, GF-2, and DEM data) to precise geopositioning of rotating scanning imagery, achieving 4.01 m planar accuracy through joint adjustment, with edge-matching RMSE between adjacent image strips down to 2.52 m.

What is striking is how consistently the fused, multi-sensor version beats the single-sensor version, often by a wide margin, and almost always on the cases (dense vegetation, cloud-covered water, spectrally ambiguous crops) where a single sensor would struggle. Declaro and Kanae (2024) put a number on this pattern for water monitoring: fusing Landsat-8/9 with Sentinel-2 and Sentinel-1 SAR shortens the median global revisit interval from roughly 16 to 23 days down to about 3.5 to 4.6 days, though harmony between the optical and radar signals drops sharply in vegetated wetlands (Kappa 0.27 to 0.45, versus 0.80 to 0.95 for open water).

3.5 Autonomous Robotics, UAVs, and Rover Navigation

The final group brings the story back to individual moving platforms that must navigate without reliable outside help.

Ušinskis et al. (2025) survey the mobile-robot sensor-fusion landscape and propose a genuinely inexpensive "channel robot" design that gets by with a camera, a couple of laser pointers, and a pseudo-LiDAR setup rather than an expensive dedicated sensor suite. Zheng et al. (2024) take the underlying philosophy to its most extreme testing ground: a Chinese desert standing in for another planet's surface. Their rover places the IMU at the center of a multi-factor graph optimization scheme that also tracks wheel ruts and the skyline for extra positioning cues, outperforming LOAM (1.59% vs. 2.18% global positioning error). It is a small but telling detail that a citrus farm in Israel, a car on a city street, and a rover in a desert all lean on the same basic idea: an IMU at the center, everything else correcting it.

4. COMPARATIVE ANALYSIS OF REVIEWED STUDIES

Table 1 summarizes the 23 reviewed studies by application domain, core technology, headline quantitative result, and principal identified gap. Figure 3 compares positioning accuracy across the navigation-focused papers on a logarithmic scale, revealing the three-order-of-magnitude spread from centimeter-level urban GNSS (Ekolama, 2026) to planetary-rover graph optimization (Zheng et al., 2024). Figure 4 presents a gap-frequency analysis showing which limitations appear most often across the literature.

Table 1. Comparative summary of the 23 reviewed studies.

S/N	Study	Application Domain	Core Technology / Method	Key Result / Principal Gap
1	Avioz et al. (2025)	Precision agriculture	UAV multispectral + Sentinel-2 + SfM, RF	Integrated R^2=0.80; CNC-yield R^2=0.66. Gap: limited cultivars; poor extrapolation
2	Susanti et al. (2026)	Mobile robot localization	GPS/IMU/LiDAR/Camera; KF/EKF/UKF/AI	Fusion cuts error 79 m→3.7 m; CKF +34%. Gap: needs adaptive, cost-effective real-time systems
3	Sharma (2026)	UAV geo-localization	Deep image matching + adaptive UKF	2.94 m RMSE; 3.5× better than VIO alone. Gap: degrades above 150 m AGL; GPUdependent
4	Wang & Ahmad (2025)	Dynamic target localization	RF/UWB/SLAM/VIO taxonomy (185 studies)	UWB 5–10 cm; SLAM most widely applied. Gap: multi-target and NLOS handling underdeveloped
5	Ekolama (2026)	Urban GNSS positioning fusion	GNN + attention multi-constellation	Median 3D error 6.2 cm; 42–57% > EKF. Gap: cold-start lag; total loss in tunnels
6	Alghamdi et al. (2025)	GPS-denied navigation	Radio/celestial/quantum nav + sensor fusion	Quantum nav sustained 4-hr flight w/o GPS. Gap: high computation and energy demand
7	Durlevic et al. (2025)	Wildfire risk mapping	MODIS/VIIRS/Sentinel-2/Landsat + DNN/XGBoost/KAN	DNN accuracy 83.4%, ROC-AUC 92.3%. Gap: coarse climate data; no temporal modelling
8	Fawakherji & Hashemi-Beni (2025)	Flood mapping	UAV RGB + Sentinel-1 SAR, RF/SVM/GTB	RF accuracy 94.12%, Kappa 91.12%. Gap: single-site validation only
9	Wang et al. (2025a)	On-orbit computing	Hyper-heterogeneous CPU/GPU/FPGA/TPU	275 TOPS/kg; >1 TOPS/W efficiency. Gap: radiation hardening and power limits
10	Wang et al. (2025b)	LEO communicationnavigation	ICAN: information/SoOP/signal augmentation	PPP convergence 30 min→<3 min with LEO. Gap: requires large constellations; complex ephemerides
11	Yuan et al. (2025)	Launch-vehicle navigation	INS+SRN+XNAV+Altimeter fused via ANN	AI: 9.1 m RMSE vs UKF 67.4 m. Gap: needs data from hundreds of trajectories

12	Hashima et al. (2025)	UAV-satellite 6G communication	DRL/MAB/FL for resource allocation	MAB gives 69–118% energy-efficiency gain. Gap: latency, spectrum and security constraints
13	Kugusheva et al. (2024)	Ocean current mapping	U-Net CNN fusing SSH, SST, chlorophyll	71.7% correct current angles (vs 67.9% AVISO). Gap: regional (Mediterranean) scope only
14	Declaro & Kanae (2024)	Surface water monitoring	Landsat-8/9 + Sentinel-2 + Sentinel-1 SAR, Otsu	Revisit interval cut from ~19 to ~4 days. Gap: weak performance in vegetated wetlands
15	Vitale & Lamonaca (2025)	Urban growth monitoring	Landsat time series + RF + transfer learning	2024 overall accuracy 92.6%; +45.9% builtup area. Gap: medium (30 m) spatial resolution
16	Dong & Yuan (2026)	Soil moisture mapping	RF/XGBoost/LightGBM/CatBoost + SHAP + Optuna	CatBoost optimal; 57% feature reduction via RFE. Gap: sparse ground stations over Tibetan Plateau
17	Al Saim & Aly (2025)	Forest/tree-species mapping	Sentinel-1/2, Landsat-8, NAIP + RF/GTB/SVM/KNN	Weighted ensemble accuracy 0.977. Gap: needs hyperspectral data for finer separation
18	Jarraya et al. (2025)	UAV GNSS-denied navigation	Absolute/relative localization taxonomy	Vision-based fusion identified as most effective. Gap: no single sensor suffices; high complexity
19	Ušinskis et al. (2025)	Mobile-robot navigation	Vision/non-vision sensor fusion architectures	Hybrid sensing improves detection & mapping. Gap: calibration and synchronization remain hard
20	Wang et al. (2025c)	Satellite image geopositioning	ZY-3/GF-2/DEM feature matching + joint adjustment	Planar accuracy 4.01 m; edge RMSE 2.52 m. Gap: sensitive to land-use change and rugged terrain
21	Guo et al. (2025)	3D semantic scene completion	Satellite-ground dual-branch deep fusion (SGFormer)	45.0 IoU / 16.7 mIoU on SemanticKITTI. Gap: sensitive to localization noise (±5 m)
22	Feng et al. (2025)	Autonomous-driving positioning	Noise-adaptive GNSS/INS fusion (ESKF)	37.7% better uncertainty modelling than baseline. Gap: confined to GNSS/INS; no LiDAR/vision layer
23	Zheng et al. (2024)	Planetary rover navigation	IMU-centred multi-factor graph optimisation	1.59% global position error vs 2.18% (LOAM). Gap: validated only in a single desert testbed

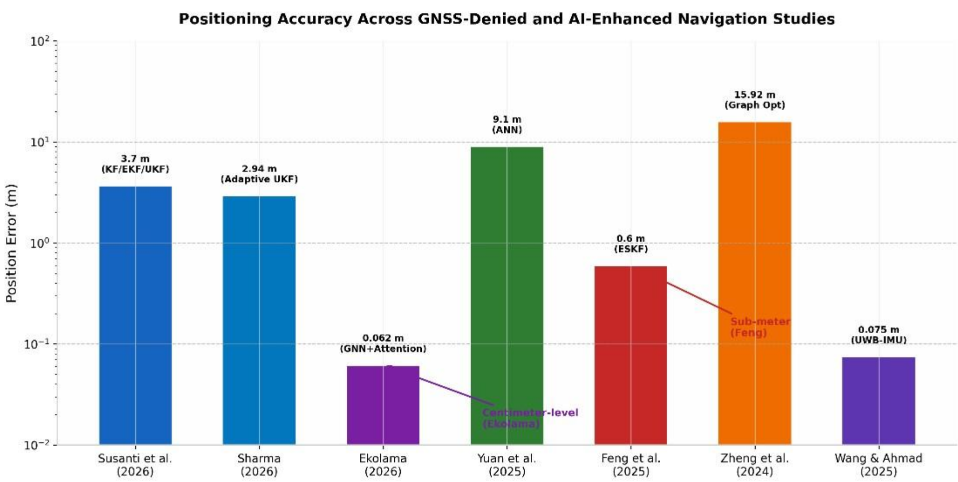

Figure 3. Positioning accuracy across GNSS-denied and AI-enhanced navigation studies (log scale).

Reading the seven bars left to right: Susanti et al.'s (2026) and Sharma's (2026) classical and adaptive Kalman-filter approaches sit in the low single-digitmeter range (3.7 m and 2.94 m). Ekolama's (2026) GNN-plus-attention fusion drops error to the centimeter level (0.062 m, the median 6.2 cm figure discussed in Section 3.2). Yuan et al.'s (2025) ANN-based launch-vehicle fusion, by contrast, sits at 9.1 m, since it operates over a fundamentally different, much longer-range navigation problem than the ground and aerial studies beside it. Feng et al.'s (2025) noise-adaptive ESKF holds autonomous-driving error to about 0.6 m, and Wang and Ahmad's (2025) survey reports UWB-IMU pedestrian localization at roughly 0.075 m, consistent with the 5 to 10 cm UWB accuracy their broader review of 185 studies identifies. Zheng et al.'s (2024) planetary-rover graph optimization is the outlier at the top of the chart, 15.92 m of absolute RMSE over an extended desert traverse, even though its 1.59% global positioning error is actually better than the LOAM baseline's
2.18%; the large absolute number simply reflects the much longer distance the rover travels between fixes, not a worse-performing method.

5. PROPOSED ARCHITECTURAL FRAMEWORK FOR FUTURE RESEARCH

Stepping back from the individual studies, a pattern becomes hard to ignore: whether the goal is to keep a drone from getting lost, map a flood, or squeeze more navigation value out of a broadband constellation, the underlying systems all seem built from the same six ingredients. That observation is the basis for the framework proposed here. It is not meant to describe any single paper's architecture exactly, since no one study implements all six layers as described, but rather to give researchers and system designers a shared reference point, a common vocabulary for the pieces they are each independently reinventing.

Figure 1 lays the framework out visually as a pipeline with a feedback loop. The subsections that follow walk through each layer, explaining not just what it does but why it exists and which reviewed studies illustrate it best.

Figure 1. Six-layer architectural framework for satellite-enabled sensing, navigation, and Earth observation, tied together by a continuous feedback loop.

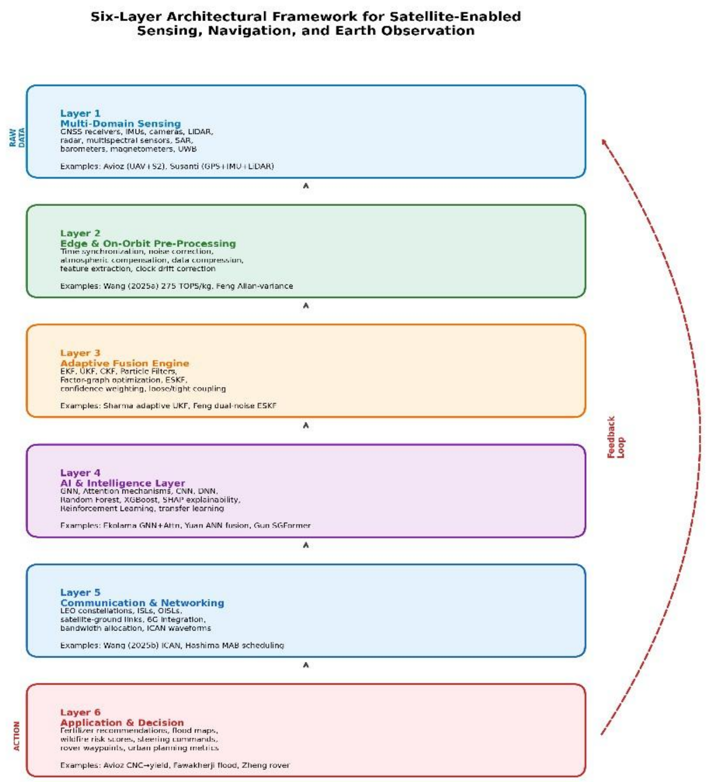

5.1 Layer 1: Multi-Domain Sensing

Everything starts here, and the defining feature is deliberate redundancy. A GNSS receiver, an IMU, a downward-facing camera, a radar unit, and a satellite passing overhead all capture something at once, each with its own blind spots. A camera cannot see through cloud; a GNSS receiver cannot see through a building; an IMU drifts if left alone. Behind almost every reviewed study, such as Avioz et al.'s (2025) combination of drone multispectral imagery with Sentinel-2 passes, or Susanti et al.'s (2026) survey of GPS, IMU, LiDAR, and camera combinations, sits the same quiet assumption: redundancy is not wasteful; it is the whole point.

The output of this layer is deliberately messy: raw, unsynchronized, arriving at wildly different rates, from a camera's 30 frames per second down to a satellite pass every few days. Untangling that mess is the job of the next layer.

5.2 Layer 2: Edge and On-Orbit Pre-Processing

Before raw sensor data is useful, it must be cleaned up and, critically, time-aligned, so that a camera frame from one instant can be compared with a GNSS fix from another. Doing this centrally, shipping every raw pixel and pseudorange to a ground station first, would be slow and would waste bandwidth. The more forward-looking systems push this work outward, onto the drone, the vehicle, or the satellite itself.

Wang et al. (2025a) document exactly this shift, reviewing on-board computing architectures that now reach efficiency levels above one TOPS per watt, making it realistic for a satellite to filter and compress its own imagery before transmission. This layer also handles the quieter but no less important task of noise correction: atmospheric distortion, sensor bias, and clock drift. Feng et al.'s (2025) Allan-variance-based noise modeling for inertial sensors fits naturally here. What emerges is a set of cleaned, time-aligned streams ready to be combined.

5.3 Layer 3: Adaptive Fusion Engine

This is the layer most people mean when they say "sensor fusion." It is where the classical mathematics of navigation engineering lives. An Extended or Unscented Kalman Filter, a particle filter, or a factor-graph optimizer takes in the cleaned streams from Layer 2 and produces a single best estimate of position, velocity, or state, weighting each sensor's contribution according to how much it should be trusted at that moment.

A GNSS fix that has just returned after a long tunnel should be trusted less, at first, than an IMU that has been running continuously. A camera trying to match features in fog should be down-weighted relative to radar. This adaptive weighting is what Sharma (2026) builds into a UKF for drone geolocalization, and what Feng et al. (2025) implement through a dual noise-estimation model for autonomous driving. Zheng et al. (2024) extend the idea to multi-factor graph optimization, where IMU pre-integration, visual reprojection, and rut-tracking residuals are jointly minimized. The output is a single, coherent estimate that carries its own uncertainty forward.

5.4 Layer 4: AI and Intelligence Layer

Classical fusion assumes noise behaves in well-mannered, Gaussian ways. Real-world sensor noise frequently does not. This layer exists to close that gap. Most reviewed studies use learned models alongside classical estimators, correcting for exactly the situations a Kalman filter handles poorly: multipath reflections off city buildings, occluded camera views, unusual satellite geometries.

Ekolama's (2026) graph neural network reasons about geometric relationships between multiple satellite constellations at once. It does not throw away the classical estimate; it improves on it, beating an EKF baseline by more than 40% in dense urban testing. Yuan, Yu, and Zong (2025) go further, letting a neural network learn the fusion rule itself. Guo et al.'s (2025) SGFormer sits in this layer too, though it corrects a camera's view rather than a position fix: a dual-branch deep network lets a ground-level camera borrow context from an overhead satellite image to fill in what its own view cannot see, lifting 3D scene-completion performance to 45.0 IoU on the SemanticKITTI benchmark, close to what LiDAR-based methods achieve. Because learned components can behave unpredictably, this layer is also where interpretability tools like SHAP earn their keep, letting an engineer ask "why did the model trust the camera less just now?" and get a real answer.

5.5 Layer 5: Communication and Networking Layer

A perfect estimate is useless if it cannot reach where it is needed in time. This layer is about the movement of fused estimates, correction data, and imagery across satellite, aerial, and ground nodes. The LEO constellation research fits here most naturally: Wang et al. (2025b) show that because LEO satellites sit so much closer to Earth, they can deliver both stronger navigation signals and a genuine communication channel, collapsing what used to be two separate infrastructures into one.

Hashima et al. (2025) address traffic management, showing how AI-based scheduling can allocate scarce bandwidth between competing UAVs and satellites far more efficiently than fixed rules. The practical effect is that a position fix computed on one platform becomes available, with minimal delay, to every other platform that needs it.

5.6 Layer 6: Application and Decision Layer

This is where all the preceding engineering turns into something a person or autonomous system can act on. A fused, intelligence-corrected estimate becomes useful once it is turned into a fertilizer recommendation for a citrus tree, a flood-risk map for a district, a steering command for a self-driving car, or a waypoint for a rover crossing unfamiliar terrain.

Nearly every application-focused study lives at this layer: Fawakherji and Hashemi-Beni's (2025) flood maps, Durlevic, Ilic, and Valjarevic's (2025) wildfire susceptibility scores, and Zheng et al.'s (2024) rover navigation commands are all decisions built on the five layers beneath them. This layer is also where outcomes and errors first become visible to humans: a farmer notices a wrong recommendation, a flood map is compared against reality, and that observation feeds the feedback loop shown in Figure 1, carrying lessons learned back up to recalibrate the sensing and fusion layers above.

5.7 Why the Feedback Loop Matters

It would be easy to draw this framework as a straight line. Many published diagrams do exactly that. But a one-way pipeline assumes that calibration, once set, stays correct forever, and none of the reviewed literature supports that assumption. Ekolama (2026) shows accuracy sagging noticeably in the first few seconds after start-up. Durlevic, Ilic, and Valjarevic (2025) note that a wildfire model trained on one season's data can drift out of step with a changing climate. The honest response is to build the loop back in: outcomes observed at Layer 6 should routinely flow up to recalibrate noise models in Layer 3 and retrain models in Layer 4. A framework without that loop is merely a snapshot of a system on its first day.

5.8 Cross-Cutting Design Considerations

Several concerns cut across all six layers and deserve explicit mention.

Robustness to signal degradation: A well-designed system should degrade gracefully, losing some accuracy rather than failing outright, when GNSS is jammed, a camera is occluded, or cloud cover blocks a satellite pass. This concern recurs across nearly every GNSS-denied navigation study reviewed here.

Computational proportionality: Not every piece of intelligence belongs on the same hardware. Lightweight inference should run at the edge, on the drone, satellite, or vehicle, while heavier model training happens on the ground or in the cloud. Wang et al. (2025a) and Sharma (2026) both designed carefully around this constraint; Wang et al. (2025a) report on-board efficiency now exceeding 275 TOPS per kilogram and better than one TOPS per watt, while Ekolama (2026), working the same edge-inference problem for GNN-based fusion, keeps prediction latency under 15 milliseconds on a Jetson AGX board.

Cold-start mitigation: Several studies (Ekolama, 2026; Durlevic et al., 2025) flag noticeably worse performance in the first moments of operation. A mature system should have a fallback estimator or sensible prior to lean on until it has gathered enough live data to trust itself fully.

Explainability and trust: As learned models take on more fusion workload, tools like SHAP need to be built into Layer 4 from the start, not bolted on afterward, especially for applications where a wrong decision has real consequences.

Standardized benchmarking: Almost every reviewed study tests its system in a different location on a different dataset using different metrics, making it genuinely difficult to say whether one approach is better than another. The field would benefit from shared, cross-domain benchmarks.

6. RESEARCH GAPS AND FUTURE DIRECTIONS

6.1 Gap-Frequency Analysis

Reading the gaps sections of all 23 papers back-to-back is a humbling exercise: the same handful of unsolved problems resurface regardless of application.
Figure 4 quantifies this. Table 2 maps each gap to the framework layers it affects, representative papers, and an assessed severity.

Table 2. Research gap matrix mapped to the proposed framework.

Gap Theme	Affected Layers	Representative Papers	Severity
Cold-start / Initialization	Layer 4	Ekolama (2026), Durlevic (2025)	High
NLOS / Multipath / Signal loss	Layer 1, Layer 3	Ekolama (2026), Feng (2025)	High
Computational cost / Power	Layer 2, Layer 4	Wang et al. (2025a), Sharma (2026)	High
Limited geographic validation	Layer 6	Avioz (2025), Zheng (2024), Fawakherji (2025)	High
Temporal misalignment	Layer 2	Avioz (2025), Declaro (2024)	Medium
Sensor calibration / sync	Layer 1, Layer 2	Ušinskis (2025), Jarraya (2025)	Medium
Real-time / latency constraints	Layer 2, Layer 5	Hashima (2025), Wang (2025a)	Medium
Explainability / trust	Layer 4	Dong & Yuan (2026), Al Saim (2025)	Medium

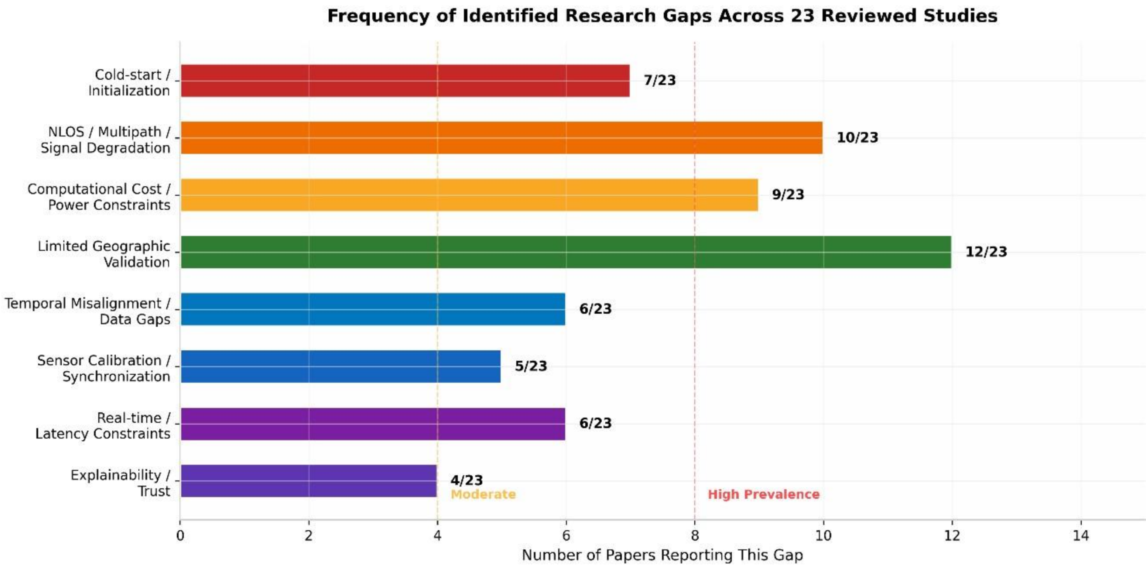

Figure 4. Frequency of identified research gaps across the 23 reviewed studies.

Figure 4 tracks eight recurring gap themes across the 23 studies. Limited geographic validation is the single most common complaint, raised in 12 of 23 papers, followed by NLOS, multipath, and signal degradation (10 of 23) and computational cost or power constraints (9 of 23), all three crossing the chart's high-prevalence line. Cold-start and initialization problems (7 of 23) sit just below that line. Temporal misalignment and real-time or latency constraints are tied at 6 of 23 each, sensor calibration and synchronization issues appear in 5 of 23, and explainability or trust concerns, the least frequently cited theme, appear in 4 of 23. Five of these eight themes are discussed in the most depth below, since they recur across the widest range of application domains rather than being confined to a single sub-field.

First is cold-start sensitivity: the awkward window right after a system switches on, before it has gathered enough live data to calibrate itself. It appears in Ekolama's (2026) urban positioning system, in deep-learning navigation more generally, and in Durlevic, Ilic, and Valjarevic's (2025) wildfire model when it encounters an unfamiliar season. None of the reviewed studies offers a fully general solution. A useful direction would be warm-start strategies that lean on a sensible prior, such as a rough map or a typical noise profile, until the system has earned the right to trust its live measurements.

Second is computational cost: a satellite, drone, or rover has a battery, and every extra layer of intelligence draws it down. Wang et al. (2025a) and Sharma (2026) both designed around this constraint. As more systems push AI-based fusion onto smaller platforms, the tension between capability and power budget will only sharpen. Model compression and hardware-aware algorithm design deserve more attention.

Third is signal degradation from non-line-of-sight conditions and multipath reflection. Both classical filters and newer learned methods still struggle in dense urban and forested settings. While attention-based models like Ekolama's (2026) show real promise, no reviewed study claims to have solved it outright.

Fourth is simply a lack of breadth in testing, the most frequently cited gap of all at 12 of 23 studies. Most studies, including Avioz et al.'s (2025) citrus orchards, Fawakherji and Hashemi-Beni's (2025) single North Carolina flood site, and Zheng et al.'s (2024) one desert testbed, validate in exactly one place. That leaves an honest question mark over how results would hold up elsewhere. Broader, multi-site validation would turn promising results into reliable ones.

A closely related concern, real-time and latency constraints, ties with temporal misalignment at 6 of 23 studies; Hashima et al. (2025) and Wang et al. (2025a) both flag it directly, since a fused estimate or a scheduling decision that arrives late is nearly as unhelpful as one that never arrives.

Fifth, and largely outside what Figure 4 quantifies, is a forward-looking concern rather than a gap any single reviewed paper reports directly: as LEO mega-constellations grow and on-orbit computing matures, questions about how thousands of satellites coordinate, how the system scales, and how it defends against cyberattack are becoming more pressing. At present, the literature addresses them only lightly, and none of the 23 studies reviewed here treats constellation-scale coordination or security as its central focus.

6.2 Technology Readiness Assessment

Table 3 estimates the maturity of the most promising technologies discussed. This assessment is based on the experimental evidence and deployment contexts reported in the reviewed studies.

Table 3. Technology readiness level (TRL) assessment.

Technology	TRL	Evidence from Reviewed Papers
LEO-based navigation (ICAN)	4–5	Wang et al. (2025b), Hashima (2025)
AI-enhanced GNSS/INS fusion	5–6	Ekolama (2026), Feng (2025)
On-orbit intelligent computing	4–5	Wang et al. (2025a)
UAV–SAR flood mapping	7–8	Fawakherji & Hashemi-Beni (2025)
ANN-based launch-vehicle navigation	3–4	Yuan et al. (2025)
Satellite-ground 3D scene completion	4–5	Guo et al. (2025)
Multi-source satellite geopositioning	6–7	Wang et al. (2025c)

7. LIMITATIONS OF THIS REVIEW

This synthesis has several limitations that readers should bear in mind. First, the temporal scope is deliberately narrow (2024–2026). That captures the leading edge of the field but necessarily omits the foundational work (classical Kalman filtering, early SLAM systems, the original GNSS constellation designs) that made the current advances possible. Second, the search was restricted to English-language sources, introducing a potential language bias. Third, the selected papers are weighted toward IEEE and MDPI journals, with fewer contributions from Springer, Elsevier, or regional journals, which may skew the methodological emphasis. Fourth, this review employs thematic synthesis rather than bibliometric network analysis; we do not claim to map citation flows or co-authorship structures. Finally, the proposed six-layer framework is a conceptual synthesis rather than a formally verified architecture; it is meant to organize thinking and highlight connections, not to replace domain-specific system designs.

8. CONCLUSION

Taken individually, these 23 studies read like reports from separate, unconnected corners of engineering. Taken together, they tell a coherent story: the future of navigation and Earth observation will not be won by any single sensor or algorithm, but by systems that combine many imperfect sources into something more trustworthy than any of them alone, and that keep learning from their own mistakes rather than treating first calibration as final.

The six-layer framework proposed here (sensing, pre-processing, fusion, intelligence, communication, and application, tied together by a feedback loop) is an attempt to give that shared story a shared vocabulary. Whether the next system is meant to fly a drone over an orchard, steer a car through a city, or guide a rover across a distant desert, the same six questions are worth asking: what is it sensing, how is it cleaning that data, how is it fusing it, what is it learning, how is it communicating, and what decision does it ultimately produce? Answering all six well, and building in the humility to keep correcting the answer over time, will matter more than any single breakthrough sensor or algorithm on its own.

REFERENCES

- Al Saim, A., & Aly, M. (2025). Enhancing tree species mapping in Arkansas' forests through machine learning and satellite data fusion: A Google Earth Engine-based approach. Journal of Geovisualization and Spatial Analysis, 9(1), 1–21.
- Alghamdi, S., Alahmari, S., Yonbawi, S., Alsaleem, K., Ateeq, F., & Almushir, F. (2025). Autonomous navigation systems in GPS-denied environments: A review of techniques and applications. In 2025 11th International Conference on Automation, Robotics, and Applications.
- Avioz, D., Linker, R., Raveh, E., Baram, S., & Paz-Kagan, T. (2025). Multi-scale remote sensing for sustainable citrus farming: Predicting canopy nitrogen content using UAV-satellite data fusion. Smart Agricultural Technology, 11, 100906.
- Declaro, A., & Kanae, S. (2024). Enhancing surface water monitoring through multi-satellite data-fusion of Landsat-8/9, Sentinel-2, and Sentinel-1 SAR. Remote Sensing, 16(17), 3329.
- Dong, Y., & Yuan, H. (2026). China's 1 km daily surface soil moisture fusion dataset (2000–2025) based on explainable machine learning. Advances in Atmospheric Sciences, 43(7), 1317–1334.
- Durlevic, U., Ilic, V., & Valjarevic, A. (2025). Wildfire susceptibility mapping using deep learning and machine learning models based on multi-sensor satellite data fusion: A case study of Serbia. Fire, 8(10), 407.
- Ekolama, S. M. (2026). AI-enhanced multi-constellation satellite fusion for centimeter-level positioning in signal-degraded environments. Journal of Optoelectronics and Communication, 8(3), 1–14.
- Fawakherji, M., & Hashemi-Beni, L. (2025). Flood detection and mapping through multi-resolution sensor fusion: Integrating UAV optical imagery and satellite SAR data. Geomatics, Natural Hazards and Risk, 16(1), 2493225.
- Feng, X., Qiu, M., Wang, T., Yao, X., Cong, H., & Zhang, Y. (2025). Noise-adaptive GNSS/INS fusion positioning for autonomous driving in complex environments. Vehicles, 7(3), 77.
- Guo, X., Hu, J., Hu, J., Bao, H., & Zhang, G. (2025). SGFormer: Satellite-ground fusion for 3D semantic scene completion. In 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Hashima, S., Gendia, A., Hatano, K., Muta, O., Nada, M. S., & Mohamed, E. M. (2025). Next-gen UAV-satellite communications: AI innovations and future prospects. IEEE Open Journal of Vehicular Technology, 6, 1998–2036.
- Jarraya, I., Al-Batati, A. S., Kadri, M. B., Abdelkader, M., Ammar, A., Boulila, W., & Koubaa, A. (2025). GNSS-denied unmanned aerial vehicle navigation: Analyzing computational complexity, sensor fusion, and localization methodologies. Satellite Navigation, 6(1), 1–32.
- Kugusheva, A., Bull, H., Moschos, E., Ioannou, A., Le Vu, B., & Stegner, A. (2024). Ocean satellite data fusion for high-resolution surface current maps. Remote Sensing, 16(7), 1182.
- Sharma, S. (2026). NGPS: GPS-denied aerial geo-localization and 2.5D reconstruction via deep satellite image matching and multi-rate sensor fusion. In 2026 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
- Susanti, V., Mahmud, M. S. A., & Saputra, R. P. (2026). Sensor fusion technology advancement in GPS-aided localization for autonomous mobile robots: A comprehensive survey. Jurnal Teknologi (Sciences & Engineering), 88(1), 165–189.
- Ušinskis, V., Nowicki, M., Dziedzickis, A., & Bučinskas, V. (2025). Sensor-fusion based navigation for autonomous mobile robot. Sensors, 25(4), 1248.
- Vitale, A., & Lamonaca, F. (2025). Advancing built-up area monitoring through multi-temporal satellite data fusion and machine learningbased geospatial analysis. Remote Sensing, 17(11), 1830.
- Wang, G., Wan, G., Su, Z., Wang, Y., Jia, Y., Li, G., & Liang, S. (2025a). High-performance on-orbit intelligent computing and real-time services for remote sensing satellites based on large-scale computing power in space. IEEE Access, 13, 12890–12909.
- Wang, L., Wang, P., Zhang, Y., Wang, Y., & Chen, B. (2025c). Multi-source collaborative positioning for ultra-wide swath rotating scanning satellite imagery based on multi-source data fusion. Sensors, 25(3), 850.
- Wang, M., Nardin, A., Ma, R., Wang, R., Dovis, F., Garello, R., & Liu, G. (2025b). Integrated communication and navigation based on LEO satellite networks: A survey. IEEE Access, 13, 46232–46258.
- Wang, S., & Ahmad, N. S. (2025). A comprehensive review on sensor fusion techniques for localization of a dynamic target in GPS-denied environments. IEEE Access, 13, 2246–2280.
- Yuan, Y., Yu, F., & Zong, H. (2025). Multisensor integrated autonomous navigation based on intelligent information fusion. Journal of Spacecraft and Rockets, 62(4), 1–11.
- Zheng, B., et al. (2024). An autonomous navigation method for planetary rover based on multi-modal fusion and multi-factor graph optimization. Journal of Physics: Conference Series, 2762, 012002.