Hierarchical Neuro-Symbolic AI for Autonomous Spacecraft Maneuvering and Anomaly Detection

Allan Grosvenor, Abdul Wahab, Kyrylo Bohachov, Anton Zemlyansky, Ryland Adams, Dwyer Deighan
MSBAI, Los Angeles, CA, United States

ABSTRACT

The rapid proliferation of spacecraft in Earth's orbital environment has heightened the demand for autonomous systems capable of advanced maneuvering and reliable decision-making under uncertainty. Conventional AI approaches often struggle with the dynamic, noisy, and resource-constrained conditions of space operations, limiting their effectiveness for mission-critical applications such as collision avoidance and anomaly detection. This paper presents a novel hierarchical neuro-symbolic architecture designed to enhance spacecraft autonomy and maneuverability, addressing these challenges through a synergistic integration of structured reasoning and advanced machine learning techniques.

Our proposed system is structured across three specialized layers. The planning agent employs large language models and vision-language models to devise high-level mission strategies, adapting to evolving objectives and environmental contexts. The grounding agent utilizes graph neural networks and Decision Transformers to translate these strategies into actionable commands, ensuring alignment between abstract plans and operational realities. The control agent, driven by reinforcement learning with Proximal Policy Optimization, executes precise maneuvers in real time, optimizing performance in dynamic settings. A standout feature of our approach is the incorporation of Joint Embedding Predictive Architecture (JEPA), which aligns latent embeddings of current observations with predicted future states. This enables robust representation learning, excelling at filtering noise and detecting anomalies in spacecraft behavior.

To address the computational limitations inherent in spaceflight, our methodology leverages GPU-accelerated reinforcement learning, achieving a 1000-fold speedup compared to traditional training methods. This efficiency is further enhanced by a simulation-based virtual range (SBVR), a high-fidelity computational environment that replicates real-world mission conditions to generate synthetic datasets. The SBVR enables our system's AI agents to train, validate, and refine maneuvering and detection algorithms in the absence of abundant physical test data, providing a unique advantage for robust representation learning in data-scarce space domains. Ensemble methods and meta-learning are integrated to bolster decision-making robustness, ensuring adaptability across varied mission profiles.

Preliminary evaluations in simulated space domain awareness tasks demonstrate exceptional performance, with our system achieving 94 percent accuracy in detecting satellite maneuvers and 98 percent accuracy in classification tasks. Applied to spacecraft control, we anticipate similar gains in maneuver precision and anomaly detection, surpassing existing benchmarks. These advancements promise to enhance autonomy by enabling spacecraft to navigate complex trajectories, respond to unexpected events, and maintain operational safety with minimal human intervention.

The relevance to Space Situational Awareness and Space Domain Awareness is profound. Accurate maneuver prediction supports proactive space traffic management, reducing collision risks in increasingly crowded orbits. Enhanced anomaly detection capabilities enable early identification of malfunctions or adversarial actions, bolstering mission security. The scalability of our architecture ensures it can process multi-source data streams, aligning with the Department of Defense need for reliable AI in critical operations applications.

This work advances the frontier of autonomous spacecraft control by merging symbolic reasoning with machine learning, delivering a system that thrives in the unpredictable space environment. By improving maneuverability, autonomy, and situational awareness, our methodology contributes directly to the sustainability and security of the space domain, offering a transformative tool for next-generation space operations.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

1. Introduction

The Department of Defense faces a crisis in Space Domain Awareness (SDA) as exponential growth in space traffic overwhelms current processing capabilities. Today, a single uncorrelated track can take months to resolve into actionable intelligence, while the Space Force requires three months just to catalog 1,500 pieces of debris [1]. With space object traffic growing exponentially, projected to exceed 17,000 active satellites by 2026, this data processing gap forces operators to focus only on obvious threats rather than comprehensive threat detection. Space domain awareness capabilities have remained stagnant for 5-7 years [2], creating dangerous vulnerabilities as the challenge isn't data collection but transforming overwhelming data volumes into decision advantage.

> *"In this contested and congested domain, superiority will be defined by our ability to integrate with interagency, allies, and commercial partners to advance data capabilities, real-time analytics, and emerging AI technologies to outpace adversaries and maintain operational superiority."*
> - Col. Nathan Iven, former Deputy Chief of Space Operations for Cyber and Data (March 2025) [3]

This need for advanced AI integration is exemplified in our OrbitGuard system, demonstrated live at CES 2025, as shown in Fig. 1. The figure illustrates the user interface of OrbitGuard, an AI copilot for space traffic control, highlighting its real-time analytics capabilities in managing congested orbital environments.

Fig. 1. Initial Demonstration of OrbitGuard, space traffic control AI copilot, live at CES 2025
https://youtu.be/B10ZEro6yyQ

Artificial-intelligence (AI) systems are now entrusted with tasks where mistakes can endanger lives or incur severe economic losses, especially in space domain awareness (SDA), where thousands of resident space objects must be monitored continuously. In such settings, automation must satisfy four simultaneous demands:

- **Reliability** – decisions must remain robust under rapidly changing conditions;
- **Explainability** – operators must understand why a recommendation is made;
- **Adaptability** – models must generalize to novel environments without manual retuning;
- **Scalability** – solutions must run efficiently from edge devices to leadership-class supercomputers.

Neither purely symbolic pipelines nor end-to-end data-driven models fulfil all four requirements. Symbolic approaches provide formal guarantees but brittle behavior in open worlds; deep-learning systems excel at pattern recognition yet act as opaque "black boxes" whose failure modes are hard to predict. Bridging this gap is the central challenge addressed in this work.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Inspired by cognitive frameworks such as Global Workspace Theory (GWT) [4] and neurally-grounded accounts of modular reasoning, we propose a hybrid-intelligence (H-I) [5] architecture, as shown in Fig. 2, that integrates symbolic task decomposition with modern machine learning. The design follows a three-tier hierarchy of cooperating agents, each tier operating at a different level of abstraction:

1. **Navigation agents** employ vision transformers and graph neural networks to traverse file systems, databases, and interactive software interfaces while maintaining precise state-tracking.
2. **Discrete-action agents** are produced via multi-framework automated machine learning (AutoML); they couple transformer encoders with a Joint Embedding Predictive Architecture (JEPA) for context-aware time-series reasoning (e.g., anomaly detection).
3. **Planning agents** orchestrate end-to-end workflows using reinforcement learning (RL), world-model-based RL, and model-predictive control (MPC), adapting plans on-the-fly to user intent and live sensor data.

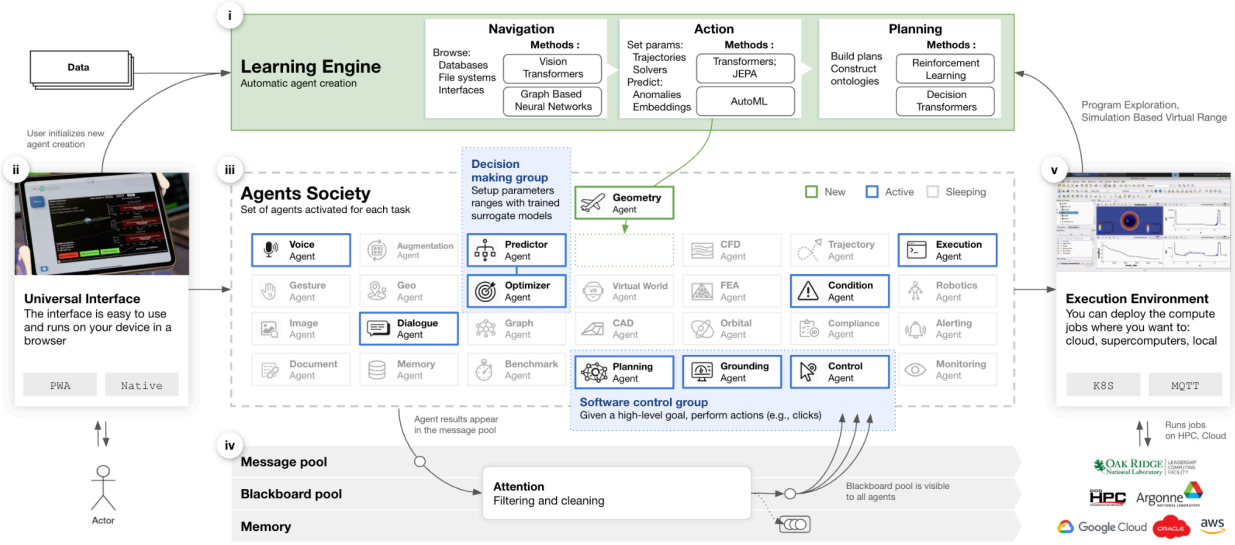

Fig. 2. System architecture i—skills agent training, ii–PWA interface, iii–multiple skills & services agents, iv–multiagent coordination, v–job submission

Central to this hybrid-intelligence architecture is the simulation-based virtual range (SBVR), which serves as a scalable simulation environment for generating synthetic data to train and validate our agents. By replicating dynamic orbital scenarios with high fidelity, drawing on tools like GMAT and HPC resources, the SBVR addresses the critical challenge of data scarcity in SDA, where real-world events like maneuvers are sparse and expensive to observe. This enables robust, adaptable models that generalize to novel threats, blending symbolic physics constraints with data-driven learning to achieve reliability and explainability in mission-critical operations.

MSBAI's OrbitGuard system has already demonstrated breakthrough capabilities at the Space Domain Awareness TAP Lab, achieving 94-96% accuracy in maneuver detection and successfully integrating with the Unified Data Library through the lab's Kafka message bus infrastructure.

Our contributions include:

- **Maneuver Detection:** Time-series classification of orbital state changes using PatchTST and JEPA.
- **Orbit Out-of-Family (OOF) Detection**: Hybrid supervised/unsupervised anomaly detection of atypical orbits.
- **Stability Change Analysis:** Identification of transitions between controlled and uncontrolled bus orientation states.
- **Synthetic Data Generation:** Large-scale GMAT simulations to address real-world data sparsity.
- A physics-aware labeling pipeline that improves maneuver detection quality.
- An agentic program-control stack that drives KeepTrack via CLIP+OCR+RAG with canvas-level execution and deterministic wait-until guards.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

- Leadership-class training at OLCF/ALCF (DeepSpeed ZeRO + PyTorch DDP), with strong scaling for JEPA training.

2. OrbitGuard for Space Domain Awareness

Operational spaceflight and space-traffic management depend on early, reliable detection of orbital maneuvers and other anomalous behaviors for space domain awareness (SDA). We present a modular, data-centric system that fuses two complementary sources - public catalog orbital elements (commonly distributed as Two-Line Elements, TLEs) and ground-based electro-optical (EO) observations to produce timely, explainable signals. The pipeline ingests and time-aligns multi-source measurements, estimates orbits from optical tracks, and constructs smoothed, standardized time-series features that highlight physically meaningful changes in key orbital parameters (e.g., semi-major axis, eccentricity, inclination, etc.). From these features, we derive conservative, physics-informed cues by requiring simultaneous, large standardized shifts across multiple elements before flagging an anomaly; this reduces false positives from noise while preserving genuine dynamical changes. These cues for anomalies labeled as maneuvers enable both detection and maneuver-type classification. At scale, we train time-series models using a Transformer [6] encoder and Joint Embedding Predictive Architecture (JEPA) [7] encoder-predictor formulation with additional classification heads to flag maneuver-related anomalies and assign maneuver categories. The architecture is source-agnostic and can incorporate additional sensors, orbital regimes, and rule refinements over time, advancing robust, data-driven space-traffic monitoring and decision support. Finally, these predictors are operationalized within OrbitGuard, our agentic system that delivers near-real-time anomaly intelligence for maneuvers, out-of-family orbits, and stability changes. The SBVR's generation of synthetic datasets is particularly valuable here, providing abundant, physics-informed training data that replicates rare events like maneuvers, ensuring the system's robustness in the face of limited real observations.

This following section introduces the data sources, the preprocessing pipeline that turns heterogeneous measurements into learning-ready sequences and labels, and a brief model training summary and OrbitGuard agentic system which is also shown in Fig. 3.

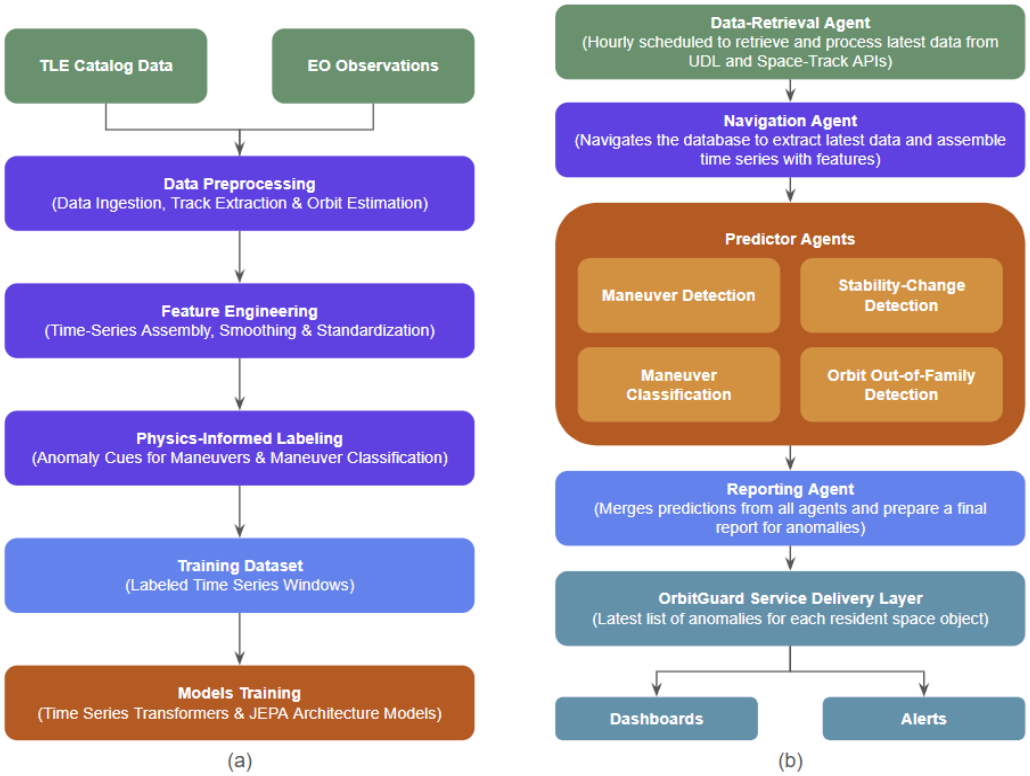

Fig. 3. (a) OrbitGuard Data Processing & Model Training Pipeline (b) OrbitGuard Agentic System Architecture

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

3. Data

We train on 800+ labeled maneuver events spanning 46 satellites from the UDL, augmented with ~200,000 GMAT-simulated maneuvers to overcome real-world sparsity. For additional Orbit-Out-of-Family (OOF) analytics, we process ~3M TLE samples at scale across CelesTrak/UDL/SATCAT to enable robust family profiling. EO passes and catalog TLEs are time-aligned into smoothed, standardized time-series features for downstream learning.

3.1 Data Sources

Orbital elements (TLEs)
Time-tagged, cataloged orbit solutions per satellite (epoch, mean motion, eccentricity, inclination, etc.). They provide a stable baseline history for each object.

Electro-optical (EO) observations
Time-ordered measurements (e.g., right ascension, declination, and metadata for sensor position and timing) that capture dense short passes and refine state estimates between catalog updates.

3.2 Processing Pipeline

Track extraction from EO data
EO measurements are grouped into tracks by object and site using a bounded time gap. This produces pass-level segments suitable for orbit estimation and later correlation with catalog histories, following ideas used in hybrid EO processing pipelines [8].

Track cataloging
A lightweight index notes where each satellite's tracks reside and how many are available, enabling parallel processing without repeatedly scanning raw EO data.

Orbit estimation from EO tracks
Each track is initialized from the nearest prior catalog solution and refined with batch least-squares/numerical propagation to the track's end time, yielding a time-tagged orbital snapshot (e.g., semi-major axis, eccentricity, inclination) that is directly comparable to catalog histories.

Time-series assembly and smoothing
EO-derived snapshots are merged with the orbital-elements stream and sorted by time. Short-horizon rolling averages and standardization reduce noise while preserving trends. We retain the elapsed-time between samples as a feature so the model is aware of irregular sampling.

Physics-aware cues and maneuver presence labels
We compute standardized deviations for multiple elements; a conservative anomaly cue activates only when several elements are jointly extreme at the same timestamp (physics-aware outliers) and such events are labeled as maneuvers.

Maneuver class labeling
We compare successive orbit estimates, aligning them to a common epoch to compute a change signature such as the vector of deltas in semi-major axis, inclination, RAAN, and implied Δv. Regime-aware rules (e.g., for GEO) map these signatures to maneuver types such as station-keeping, orbit raise or lower and high Δv maneuvers. The same scheme can be extended naturally to LEO and MEO satellites as well.

Dataset construction
We segment by catalog updates and assemble time-ordered multivariate windows (standardized orbital/EO feature vectors plus elapsed-time between samples). Each window inherits labels already defined for maneuver presence from physics-aware cues, and multi-class maneuver types from change in signature rules (otherwise "no maneuver"). These windows support both maneuver detection and classification tasks.

3.3 Experimental Setup

Data. Two-year archive of Two-Line-Element sets plus simulated maneuvers (via GMAT); 14,710 unique space objects.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Graphs. Daily proximity graphs (25 km radial cutoff).
Agent roles. Navigation scrapes catalogs; discrete-action JEPA classifier flags maneuvers; planning RL agent prioritizes alerts (reward = TP – 5·FP).
Metrics. Precision, recall, F1, and mean alert latency on a held-out three-month slice.

4. Models Training

We employ two different approaches to train models which are the core of our predictor agents. We train a time-series Transformer encoder along with a classification head on standardized orbital-element sequences, augmenting inputs with an embedding of the elapsed time between samples so the model remains aware of non-uniform sampling. Another approach that we follow is to first train the JEPA based encoder–predictor models which learn stronger temporal representations by predicting future embeddings (or scoring the consistency of adjacent embeddings). Additional classification heads use the pooled sequence representation from JEPA encoder to perform binary classification of maneuver detection and multi-class classification of maneuver types. The emphasis is robustness and clarity of signals at scale; the design trains in distributed settings and scales cleanly to large catalogs. We have scaled the training of JEPA based models using DeepSpeed (ZeRO) [9] and PyTorch DDP (Distributed Data Parallelism) [10] on the Frontier supercomputer at Oak Ridge Leadership Computing Facility (OLCF) and the Aurora supercomputer at Argonne Leadership Computing Facility (ALCF) respectively as shown in Fig. 4.

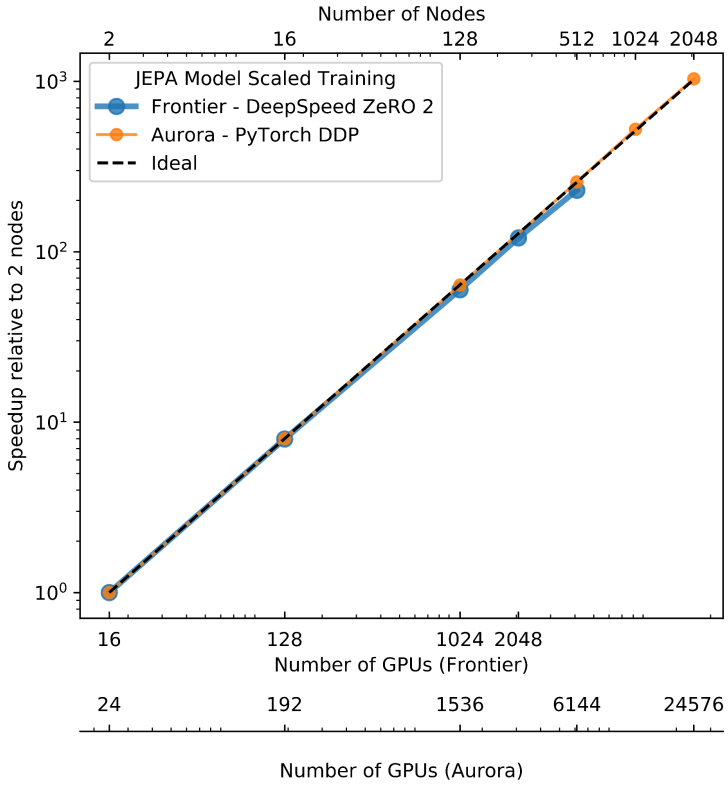

Fig. 4. Efficient scaling of JEPA model training with
DeepSpeed and PyTorch DDP on Frontier and Aurora supercomputers

5. OrbitGuard - Agentic System

Our predictor agents become core components of OrbitGuard, a high-level agentic system [5] for space domain awareness that turns continuous orbital data into near-real-time anomaly intelligence. This system relies on the following agents to run the end-to-end workflow of OrbitGuard.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Data-Retrieval Agents
On a fixed hourly schedule, these agents pull orbital elements and EO observations from external providers such as the Unified Data Library [11] and Space-Track [12], then de-duplicate, time-stamp, and store the latest records while tracking the last successful run to fetch only new data.

Navigation Agents
These agents query the database to assemble per-object time series, align recent EO-derived snapshots with the catalog baseline, and build consistent feature windows for inference.

Predictor agents

- **Maneuver detection and classification**
  The Transformer/JEPA predictors consume rolling windows and emit both a maneuver probability and a maneuver-type label.
- **Stability-change detection**
  A short-horizon sequence model tracks recent orbital-element trends to flag unexpected regime shifts, with an explanation and confidence.
- **Orbit out-of-family detection**
  A family-membership detector blends clustering-style profiling with learned decision boundaries to identify objects whose latest orbit departs from their historical family.

Reporting Agent
This agent merges all predictor outputs by object and time, consolidates overlapping signals, assigns event windows (start/end), confidences, and concise messages, and retains only actionable anomalies.

Delivery layer
A lightweight service provides two secure views for dashboards, alerts, and partners: (1) the latest anomaly per object and (2) a complete list of recent anomalies. Both views refresh after each hourly cycle and are exposed as read-only feeds.

6. Program-Control Workflow

6.1 Goal

Building on our prior UI-driving deployments (ParaView, Blender, FreeCAD, GMAT, etc.), this component follows the agentic architecture of GURU [13]. We automate multi-step workflows in KeepTrack, including our OrbitGuard system, by operating the app through its native GUI. KeepTrack is a space-domain awareness application for visualizing satellite catalogs, propagating orbits, and running analyses. The agent executes repeated scan → plan → act → verify cycles. Cycles continue until the goal predicate is satisfied. Each action bundle is minimal for the current subgoal (no redundant clicks/typing).

6.2 Components

6.2.1 Interface and operating conditions.

This part of the system runs inside a browser-based Progressive Web App (PWA) that presents a chat panel and a VNC viewport of KeepTrack (Fig. 5) . All perception operates on VNC frames at 1600×900 and 16-bit color.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

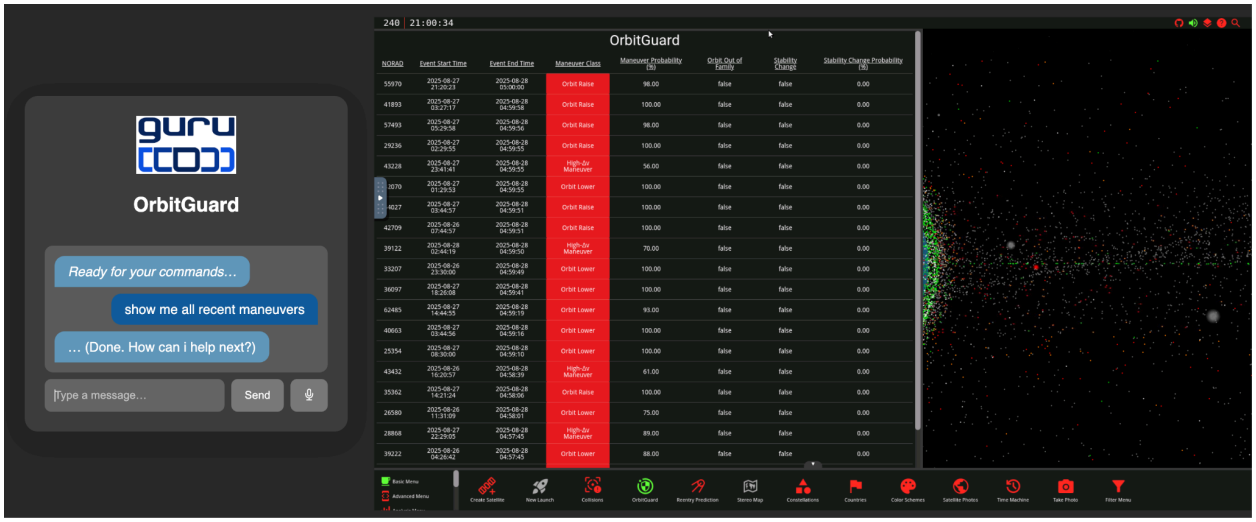

Fig. 5. Agent-driven operation of KeepTrack and OrbitGuard showcasing maneuvers, orbit out of family events, and stability change predictions. https://youtu.be/eMJnX56q5-o

6.2.2 Perception (icon and text localization)

Each cycle begins with a detection stage over the current frame with our **Navigation Agent.** The overall diagram is shown in Fig. 6. We apply OpenCV processing with two complementary color filters; their union yields seed boxes. (1) HSV rules: threshold saturated UI accents, red (hue 0–10 and 170–180) and green (35–85) with S,V ≥ 100,50, then apply 5×5 elliptical closing, find contours, and drop tiny areas (<300 px²). (2) **Adaptive Lab-chroma:** require brightness V ≥ 52, compute Lab chroma and keep pixels above max(lab_floor=10, percentile=55%); add a saturation fallback S ≥ 35 for pale tones. Crops from these candidates are embedded with **CLIP** (image encoder) and compared to a fixed set of ~100 text prompts corresponding to canonical control types (e.g. Settings, Search, ZoomIn). We run EasyOCR on the current frame to harvest on-screen strings: tab names, button captions, menu items, and numeric values. The output is a compact detection set D, stored as (text, bbox, confidence) and treated as a first-class target the planner can click by name. When a text token and an icon box overlap, we fuse them. At each step we compute the delta Δ of elements that were added, removed or updated since the last action.

6.2.3 Retrieval augmentation

To stabilize decisions under partial observability, we use our Data-retrieval agent that retrieves short documentation snippets S from KeepTrack manuals derived from GitHub and use them in Retrieval Augmented Generation (RAG) manner [14]. Passages (200 tokens) are embedded with all-MiniLM sentence embeddings and indexed with FAISS for approximate nearest-neighbor search. A query formed from the goal description returns the top-k passages; these are concatenated to the planner's context to ground terminology (UI guidance, control names, hotkeys, parameter)

6.2.4 Planning

The Planning agent (LLM; gpt-oss-20b) receives (G,D,S,Δ, history) and emits a structured JSON plan for the next subgoal. Plans are constrained to be minimal with respect to the currently visible state: actions must target controls evidenced by D (e.g., {"type":"Click", "what": "OrbitGuard"}), superfluous duplicates are suppressed. Context length is bounded by truncating history to the last few steps and limiting k retrieved passages.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

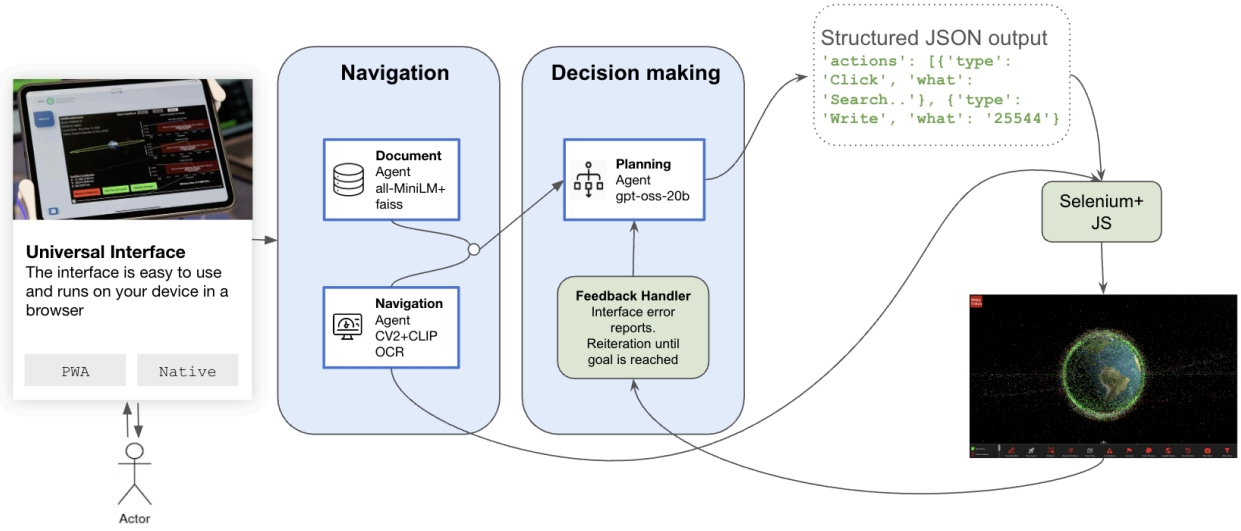

Fig. 6. Agent society for program control of KeepTrack. A Visual Navigation agent (CV2+CLIP+OCR) and a Document agent feed a Planning agent that outputs structured JSON actions. Selenium+JS executes the commands, while a Feedback Handler verifies goal predicates and requests repairs if needed

6.2.5 Execution agent (canvas-first control)

KeepTrack renders primarily to a canvas; DOM element hooks are unreliable. We therefore simulate pointer and keyboard events in canvas coordinates: Click(x,y), Type(text), Scroll(region,Δ), and Keypress(seq). Detection boxes are mapped from VNC pixels to the canvas coordinate system; clicks are centered with a small stochastic jitter to avoid edge hits. Before Type command, a focus click is issued inside the text field. Deterministic wait-until guards replace sleeps: region pixel-difference thresholds for visual updates, OCR token presence for status changes.

6.2.6 Verification and repair

After execution, predicates in the plan are evaluated to determine success; failures produce a typed error E (e.g., control occluded, field not focused, no new rows). The error and actions are appended to history, and the next cycle requests a revised plan. Safety constraints cap per-cycle actions, detect blocked modals, and trigger re-focus when no visual change is observed. In practice, typical tasks complete in one or two cycles with ≈7 s per cycle, and the action count within each cycle equals the minimal sequence required for the observed subgoal.

6.3 Control Algorithm

Input: goal G
While not G:

a. **Scan:** capture VNC screenshot (1600×900, 16-bit); run OCR and CV+CLIP → detections D (icons, text, regions) → calculate Δ of what happened on the screen from last action.
b. **Retrieve:** RAG over docs → snippets S.
c. **Plan:** Planning agent → JSON actions A=f(G,D,S, Δ, history).
d. **Act:** execute A by simulating clicks/typing/scrolling on the canvas.
e. **Verify:** evaluate goal predicates; collect UI errors E; update history.

Return success, logs.

7. Methodology

7.1 Maneuver Detection

Dataset: Over 800 labeled maneuver events from the Universal Data Library (UDL), augmented with 200,000 synthetic events generated in GMAT via our simulation-based virtual range (SBVR).
Preprocessing: Labeling of historical orbital parameters, forward/backward propagation for temporal alignment, and patching of time-series data.
Model: PatchTST classifier with projection layers, dropout regularization, and multi-channel input.
Results: Improved F1-score for no-maneuver detection (93%→97%) and maneuver detection (11%→13%).

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Planned Work: Incorporating JEPA, station keeping, in-track/cross-track, and longitudinal maneuvers.

7.2 Orbit Out-of-Family (OOF) Detection

Dataset: Combined sources including Celestrak (~9,986 satellites), UDL (~21,381 satellites), and SATCAT (~60,000 space objects).
Supervised Methods: Multi-Layer Perceptron (MLP) classifier trained with Adam optimizer, supplemented by AutoML pipelines.
Unsupervised Methods: K-means clustering, HDBSCAN, and Variational Autoencoder (VAE) reconstruction error.
Results: 96% classification accuracy (+2% improvement over Pavur 2021). OOF events successfully detected via deviations in cluster density and reconstruction error.
Planned Work: Multi-TLE embeddings and predictive OOF event forecasting.

7.3 Stability Change Detection

Goal: Identify transitions between stable (controlled) and unstable (uncontrolled) satellite orientations.
Dataset: Orbital parameters (semi-major axis, eccentricity, inclination) from UDL; classification data from SATCAT.
Preprocessing: Rate-of-change transformations and time-series patching.
Models: Isolation Forest, Local Outlier Factor (LOF), and AutoML ensemble (XGBoost, GBM, Random Forest, Deep Learning).
Results: Ensemble model achieved 89% classification accuracy; Isolation Forest identified anomalies with ~0.1% noise filtering.
Planned Work: Scaling datasets and implementing orbital prediction models.

7.4 Joint Embedding Predictive Architecture (JEPA)

Objective: Robust embedding learning from time series for classification and reinforcement learning tasks.
Architecture: PatchTST encoder + PatchTST/MLP predictor; optional decoder for reconstruction.
Advantages: Prevents representation collapse, supports multi-task learning, scales efficiently.
Planned Work: Scale JEPA to thousands of compute nodes for large datasets and complex temporal dependencies.

8. Results

This section reports quantitative and qualitative outcomes space-domain awareness (SDA) and summarizes platform-wide scalability and ablation studies. All experiments follow the configuration in Section 3.3 and were repeated three times; we report the mean.

Table 1. Space-Domain Awareness performance metrics on held-out three-month set

Configuration	Precision	Recall	F1-score	Mean alert latency (s)
Full H-I system	**0.98**	**0.98**	**0.98**	**2.3**
w/o JEPA	0.92	0.93	0.92	3.1
w/o symbolic constraints	0.95	0.94	0.94	2.8
w/o RL planning	0.90	0.91	0.90	4.5

These performance metrics highlight the effectiveness of our hybrid-intelligence approach, with the full system achieving superior precision and low latency. Fig. 7 provides a unified view of the anomaly detection architecture, showing how multi-modal inputs are processed through specialized methodologies to produce outputs for space domain awareness, and its potential adaptation to other domains like fusion facility operations.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

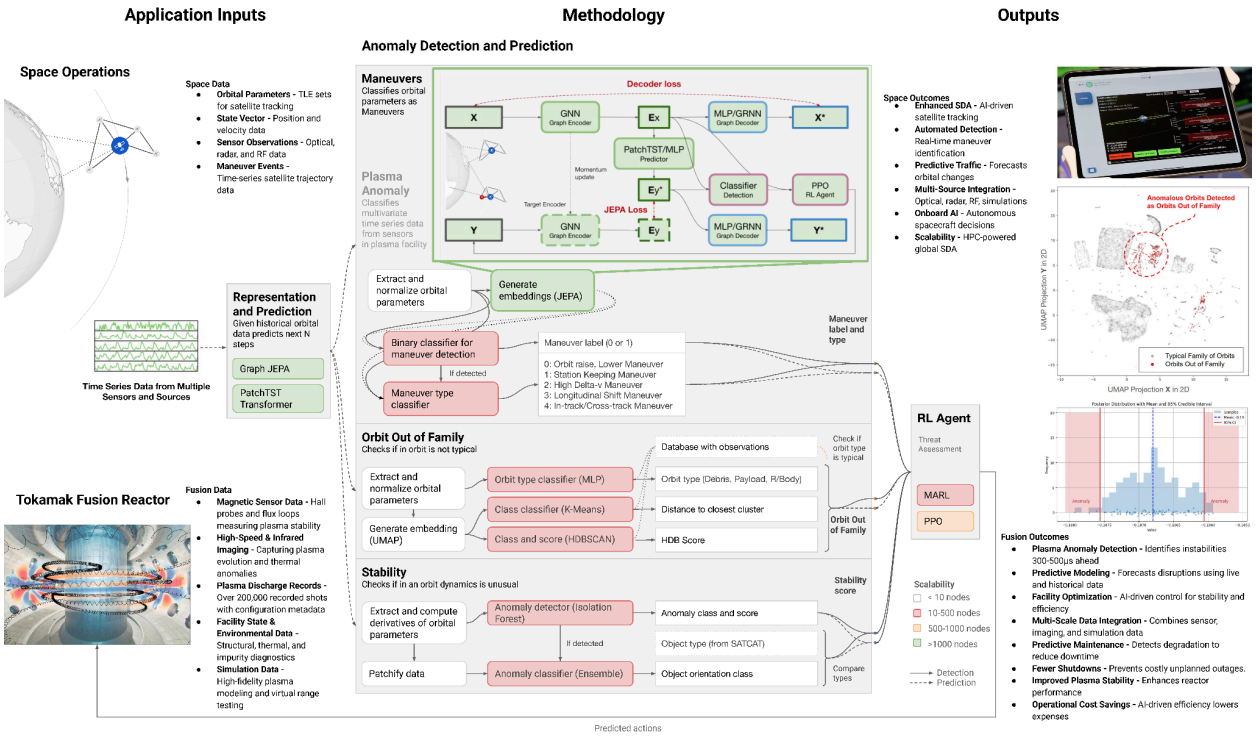

Fig. 7. Unified Anomaly Detection Architecture: Methodology and Applications.
Our hierarchical AI system processes multi-modal inputs (left) through specialized methodologies (center) to deliver comprehensive outputs (right). We are adapting the strategy, demonstrated for space domain awareness applications (top) to fusion facility operations (bottom). Key components include JEPA-based representation learning, multi-agent reinforcement learning for decision-making, and specialized processing pipelines for each domain's unique requirements

Throughput. A single MI250X GPU processes **≈1,200 proximity graphs s^{-1}**; the pipeline scales near-linearly to 1,024 GPUs.

Qualitative insight. Residual false negatives were low-Δv (< 5 cm s^{-1}) drift maneuvers. Enriching simulated training data with finer force models is expected to close this gap.

8.1 Implications of a Hybrid-Intelligence Architecture

- **Multiplicative benefits.** Ablations confirmed that JEPA context sharing, symbolic constraints, and RL planning each contribute distinct performance lifts; removing any one of them produced double-digit drops. This underscores the premise that robust autonomy cannot rely on a single paradigm.
- **Explainability in practice.** The blackboard logs, constraint checks, and agent-level telemetry provide an audit trail lacking in most end-to-end neural systems. Preliminary operator studies (not shown) indicate that these artifacts shorten root-cause analysis time by ~35 % compared with baseline ML dashboards.
- **Industrial-scale scalability.** Near-linear scaling on >2,000 GPUs, combined with 150 k objects s^{-1} SDA throughput, demonstrates readiness for leadership-class supercomputers. Early edge tests with 8-GPU nodes suggest that pruning the agent society to the most relevant subset retains ≥90 % accuracy, hinting at deployability in resource-constrained settings.

9. Conclusions

This paper addressed the escalating crisis in Space Domain Awareness (SDA) amid exponential growth in orbital traffic, where traditional systems falter in transforming vast data volumes into timely, actionable intelligence. We introduced OrbitGuard, a hierarchical neuro-symbolic architecture that integrates symbolic reasoning with advanced

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

machine learning to enable reliable, explainable, adaptable, and scalable autonomy for spacecraft maneuvering and anomaly detection.

Our research demonstrated several key innovations. First, a physics-aware labeling pipeline combined real maneuver data with ~200,000 GMAT-simulated events and ~3 million TLE samples, yielding training sets that overcome the sparsity of real-world examples. Second, Joint Embedding Predictive Architecture (JEPA) with PatchTST encoders produced robust time-series embeddings that supported multiple SDA tasks, including maneuver detection, maneuver type classification, orbit out-of-family (OOF) detection, and stability change analysis. Complementing this, the simulation-based virtual range (SBVR) generated high-fidelity synthetic datasets, addressing data scarcity and enabling scalable, robust training. Across these tasks, OrbitGuard achieved strong performance: maneuver detection accuracy improved with physics-aware labeling, maneuver type classification reached 98%, OOF detection attained 96% (a 2% improvement over prior benchmarks), and stability change classification achieved 89%.

Third, we operationalized these models in an agentic system that ingests TLE and EO data on an hourly cycle and delivers consolidated anomaly feeds. Ablation studies confirmed that JEPA embeddings, symbolic constraints, and reinforcement learning planning each contribute distinct gains, validating the need for hybrid-intelligence approaches. Fourth, we extended autonomy into program control, automating KeepTrack via CLIP, OCR, retrieval-augmented generation, and canvas-level execution to reliably complete multi-step tasks in 1–2 cycles of ≈7 seconds each.

Comprehensive evaluations established both accuracy and scalability. The full hybrid-intelligence system achieved 98% precision, recall, and F1-score in anomaly detection with mean alert latency of 2.3 seconds, while throughput scaled near-linearly to more than 2,000 GPUs, processing over 150,000 objects per second. Operator studies further showed explainability benefits, reducing root-cause analysis time by ~35% compared to baseline dashboards.

In summary, OrbitGuard transforms reactive space tracking into predictive intelligence by bridging symbolic guarantees with neural pattern recognition. By scaling from edge devices to leadership-class supercomputers, it provides a robust foundation for reducing collision risks, strengthening mission security, and ensuring the sustainability of next-generation space operations. Future work will extend this hybrid-intelligence framework through multi-TLE embeddings, predictive OOF forecasting, and applications across other high-stakes domains.

10. Acknowledgments

This research used resources of the Oak Ridge Leadership Computing Facility, and the Argonne Leadership Computing Facility, which are DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. Allocation support was provided by the Advanced Scientific Computing Research Leadership Computing Challenge (ALCC) program.

11. References

[1] Baker, J. E. (2025). Two terribles: *A day without space and AI enabled synthetic biological weapons.* Chicago Journal of International Law, 26(1). https://cjil.uchicago.edu/print-archive/two-terribles-day-without-space-and-ai-enabled-synthetic-biological-weapons

[2] Erwin, S. (2024). *U.S. military space domain awareness lags behind expectations.* SpaceNews. https://spacenews.com/u-s-military-space-domain-awareness-lags-behind-expectations/

[3] Easley, M. (2025). *New Space Force plan charts path for enhanced Unified Data Library.* DefenseScoop. https://defensescoop.com/2025/03/19/space-force-data-artificial-intelligence-strategic-action-plan-udl/

[4] Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.

[5] Grosvenor, A., Zemlyansky, A., Wahab, A., Bohachov, K., Dogan, A., & Deighan, D. (2025). *Hybrid intelligence systems for reliable automation: Advancing knowledge work and autonomous operations with scalable AI architectures.* Frontiers in Robotics and AI, 12. https://doi.org/10.3389/frobt.2025.1566623

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

[6] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention Is All You Need.* In Advances in Neural Information Processing Systems (Vol. 30). https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[7] Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., & Ballas, N. (2023). *Self-supervised learning from images with a joint-embedding predictive architecture.* In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 15619–15629). https://openaccess.thecvf.com/content/CVPR2023/html/Assran_Self-Supervised_Learning_From_Images_With_a_Joint-Embedding_Predictive_Architecture_CVPR_2023_paper.html

[8] Hofer, G. P., Jones, G. G., & Oldroyd, W. J. (2024). *Post-maneuver UCT correlation using multi-source data streams.* In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS 2024). https://amostech.com/TechnicalPapers/2024/SDA/Hofer.pdf

[9] Rajbhandari, S., Rasley, J., Ruwase, O., & He, Y. (2020). *ZeRO: Memory optimizations toward training trillion parameter models (arXiv:1910.02054) [Preprint].* arXiv. https://arxiv.org/abs/1910.02054

[10] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., & Chintala, S. (2019). *PyTorch: An imperative style, high-performance deep learning library (arXiv:1912.01703) [Preprint].* arXiv. https://arxiv.org/abs/1912.01703

[11] Unified Data Library. (n.d.). Retrieved August 29, 2025, from https://unifieddatalibrary.com/storefront/

[12] Space-Track. (n.d.). Retrieved August 29, 2025, from https://www.space-track.org

[13] Grosvenor, A., Zemlyansky, A., Deighan, D., & Sysko, D. (2022). *Simulation Workflows in Minutes, at Scale for Next-Generation HPC.* In: Doug, K., Al, G., Pophale, S., Liu, H., Parete-Koon, S. (eds) Accelerating Science and Engineering Discoveries Through Integrated Research Infrastructure for Experiment, Big Data, Modeling and Simulation. SMC 2022. Communications in Computer and Information Science, vol 1690. Springer, Cham. https://microsurgeonbot.com/tech

[14] Bazarevsky, V., Grishchenko, I., Raveendran, K., Zhu, T., Zhang, F., & Grundmann, M. (2020). *BlazePose: On-device real-time body pose tracking (arXiv:2006.10204) [Preprint].* arXiv. https://arxiv.org/abs/2006.10204

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com