SpacSec 2026 Research Brief



Comparative Analysis of Machine Learning Algorithms for Anomaly Detection in Time Series of Constrained Node Operations

Authors: S.S. Sokolov, D.S. Mahov, V.P. Kochin

Affiliations: Admiral Makarov State University, Special Technology Center, Belarusian State University

Executive Thesis

Objective: Balancing anomaly detection F1 accuracy with sub-second compute latency for real-time constrained environments.

Core Finding: Ensemble spatial separation (Isolation Forest) delivers optimal real-time resilience for SWaP-limited **operational nodes**, matching the statistical accuracy of deep autoregressive models (CatBoost, LSTM-AE) at a fraction of the computational footprint.

Orbital Compute & SWaP Constraints

Parameter: Deployment Environment
Value: Distributed operational nodes with strict Size, Weight, and Power (SWaP) limits.

Challenge:
- Advanced neural models demand continuous GPU provisioning and prolonged calibration.
- Incompatible with edge-node processing budgets (CPU-bound, < 2 GB RAM).

Data Proxy:
- Analyzed 1.23M continuous telemetry operations over 30 days.
- Mimics dense, incident-free operational baselines requiring localized inference.

Cyber Warfare & EW Threat Vectors

Parameter: Threat Environment
Value: Continuous, unlabeled, and heterogeneous multidimensional data streams.

Vector 1: Point Anomalies
- High-amplitude single activity spikes.
- Proxy for instantaneous RF injection or compromised commands.

Vector 2: Collective Anomalies
- Sustained pattern changes over extended time intervals.
- Proxy for gradual EW jamming or covert exfiltration.

Limitation:
- Traditional static thresholds trigger high False Positive Rates (FPR) under legitimate seasonal background drift.

SOTA Diagnostic Limitations

- **Static SIEM Rules:**
  - Zero adaptation to baseline drift (High FPR).
- **ARIMA / Z-Score:**
  - Blind to nonlinear inter-feature dependencies.
- **Deep Autoencoders:**
  - Mandate long pristine histories and heavy SWaP allocation.
  - Structurally unsuited for immediate real-time uncalibrated deployment.

Formal Problem Formulation

$$\text{System State Vector: } \mathbf{x}_t = [x_{t,1}, \ldots, x_{t,8}]^T \in \mathbb{R}^8$$

Definition: A multidimensional, weakly-labeled time series requiring dynamic isolation of $\mathbf{x}_t$ without pre-existing ground truth.

Target Optimization Criteria

Objective Function: Maximize detection F1-score for uncharacterized vectors while strictly minimizing computational complexity O(N), bounding memory usage, and eliminating chronological calibration leakage.

Architecture & Data Topology

Parameter: Vector Generation

Value: 82 unique node identifiers (pseudonymized via SHA-256).

Parameter: Stream Composition

Value: 1,238,412 total operations
- Read: 62%
- Write: 21%
- Modify: 12%
- Delete: 5%

Parameter: Aggregation Latency

Value: 1-minute bins forming an 8-dimensional operational vector. Empty bins retained as zero-state.

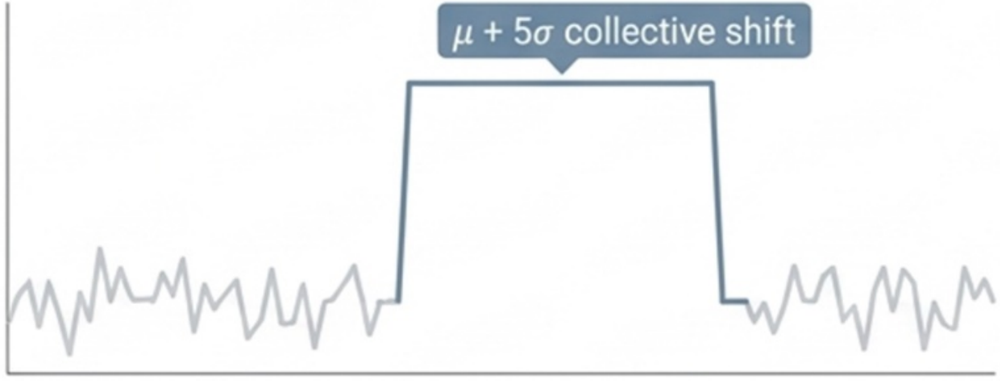
Adversary Profile & Injection Bounds

Point Injection Vector: $x'_{t_0} = x_{t_0} + \alpha\sigma$

Parameter: $\alpha \in \{3, 5, 7\}$, isolated to a single 1-minute operational bin.

Collective Injection Vector: $x'_t = \mu + \alpha\sigma$

Parameter: Maintained over sustained window $T \in \{15, 30, 60\}$ minutes.



Spatial Separation | Isolation Forest

Mechanism: Isolates anomalies in randomly constructed trees.

Governing Equation:

$$s_{IF}(x_t) = 2^{-\frac{E(h(x_t))}{c(\psi)}}$$

Where: h = path length, $c(\psi)$ = expected path-length normalization for subsample size ψ.

Operational Trait: Linear complexity, ignores sequential lag by default.

Autoregressive | CatBoost

Mechanism: Forecasts expected behavior on enriched temporal features.

Governing Equation:

$$s_{CB}(x_t) = |x_t - \hat{x}_{t+1}|$$

Feature Space: $z_t^{CB} \in \mathbb{R}^{68}$ (utilizing strict historical lag windows).

Operational Trait: High reliance on rolling statistics and chronological separation.

Reconstruction | LSTM-AE

Mechanism: Compresses and reproduces sequential normal patterns.

Governing Equation:

$$s_{AE}(x_t) = \frac{1}{8}\sum_{k=1}^{8}(x_{t,k} - \hat{x}_{t,k})^2$$

Architecture: 8-12-6-3-6-12-8 (3-unit bottleneck).

Operational Trait: Highest architectural complexity, relies on Mean Squared Error as anomaly score.

Hardware Architecture & SWaP Bounds

CPU Baseline (IF/CatBoost)	GPU Constraint (LSTM-AE)	Storage & OS
Spec: Intel Xeon Silver 4314 **Config:** 8 vCPUs allocated **Limit:** 32 GB RAM limit	**Spec:** 1x NVIDIA T4 **VRAM:** 16 GB **Driver:** CUDA 12.1 / cuDNN 8.9.2	**Disk:** 500 GB SSD **OS:** Ubuntu 22.04 LTS

Software Pipeline & Optimization

Thresholding Logic: Calibration frozen at 95th percentile prior (Q_0.95); operational limit selected via $\text{argmax}_{\tau}\, F1_{validation}(\tau)$.

Hyperparameter Tuning: Optuna TPE. 150 trials for CatBoost, minimizing validation WAPE from 39.48% to 3.80%.

Score Normalization: $s'_{m,t} = \frac{s_{m,t} - s_m^{\min}}{s_m^{\max} - s_m^{\min}}$ (frozen extrema mapping).

Dataset Parameters

Total Duration: 30 consecutive days.

Chronological Split:
-> 21 Training Days (731k events)
-> 7 Validation Days (246k events)
-> 2 Test Days (260k events)

Target Sample: 2,880 one-minute test samples.

Anomaly Injection Protocol

Point Threats:
30 isolated minutes (Distributed at 3σ, 5σ, 7σ).

Collective Threats:
15 intervals spanning 525 minutes (Lengths: 15, 30, 60 mins).

Imbalance Ratio (IR):
2325 / 555 = 4.189 (80.7% normal baseline).

Evaluation Metrics

Primary Vectors:
Minute-level F1 Score, Recall, Precision.

Secondary (Threshold-Independent):
PR-AUC, ROC-AUC.

Constraint:
Simple accuracy discarded due to always-normal baseline achieving 80.73% automatically.

Overall F1 Score Matrix

CatBoost: 0.920 F1 (Highest Point F1: 0.915)

LSTM-AE: 0.910 F1 (Highest Collective F1: 0.903)

Isolation Forest: 0.893 F1

Computational Overhead & Latency

CatBoost: 420 s train | 9 ms infer | 4 CPU cores | 3.2 GB RAM

LSTM-AE: 2,640 s train | 21 ms infer | 1 NVIDIA T4 | 5.8 GB VRAM

Isolation Forest: 8 s train | 12 ms infer | 1 CPU core | 0.9 GB RAM (Optimal SWaP)

Statistical Significance

- Method: 5-fold exact two-sided Wilcoxon signed-rank test.
- Result: p >= 0.0625 across all pairings.
- Finding: Differences fail the Bonferroni-adjusted significance level (0.0167). Superior accuracy of deep models is not statistically proven.

RF/Injection Robustness

ROC-AUC Performance: CatBoost (0.963) -> LSTM-AE (0.956) -> IF (0.944).

PR-AUC Resiliency: Ranges from 0.913 to 0.944, yielding a 4.7x to 4.9x performance multiplier over the random baseline (0.1927).

Threat Detection Delay | MTTR Proxy

Point Event Delay: Bounded to <= 1 minute across all models.

Collective Event Delay: CatBoost & LSTM-AE: 2 min median latency.

Isolation Forest: 3 min median latency.

Concept Drift Monitoring

Vulnerability: All unsupervised bases degrade under legitimate operational shifts.

Kullback-Leibler Tracking:

$$D_{KL}(P_t \parallel P_{ref}) = \sum P_t(i) \ln \frac{P_t(i)}{P_{ref}(i)})$$

Trigger: Recalibration mandated when $D_{KL} > 0.10$ for three consecutive daily windows.

SpacSec 2026 Deployment Diagnostic Matrix

Metric	Isolation Forest	CatBoost	LSTM Autoencoder
Optimal Use Case	Real-time constrained nodes	Retrospective ground verification	Collective EW pattern detection
Overall F1 Score	0.893	**0.920**	0.910
Statistically Superior?	Baseline	No (p=0.0625)	No (p=0.0625)
Peak Memory Footprint	**0.9 GB RAM**	3.2 GB RAM	5.8 GB VRAM
Hardware Requirement	**1 CPU Core**	4 CPU Cores	NVIDIA T4 GPU
Training Calibration Cost	**8 seconds**	420 seconds	2,640 seconds
Collective Alert Delay	3 minutes	**2 minutes**	**2 minutes**

Model Boundary Conditions

Limitation: Slow Exfiltration
-> Fails to detect coordinated data leakage below the 3σ amplitude threshold.

Limitation: Dispersed Activity
-> Blind to coordinated multi-user swarm behavior (features are presently calculated per node/user).

Limitation: Rare ACL Changes
-> Aggregate vector channels drown out sparse, high-privilege configuration changes.

Space-Grade Mitigations

Mitigation 1:
-> Deploy parallel 7-day and 30-day cumulative volume detectors strictly for the sub-3σ exfiltration gap.

Mitigation 2:
-> Integrate cohort and graph-based network correlation features for swarm detection.

Mitigation 3:
-> Apply higher algorithmic weights to rare-event operations to prevent aggregate masking.

Validated SpacSec Directives

- **Finding 1 (The SWaP Verdict):** Isolation Forest is the mandated choice for real-time edge nodes. It requires 86x less RAM and 330x less training time than LSTM-AE, without a statistically significant drop in F1 detection capability.
- **Finding 2 (Ground-Station Retrospective):** CatBoost provides the highest absolute point-event F1 (0.915) and is recommended for unconstrained, retrospective telemetry verification.
- **Finding 3 (Collective Resilience):** LSTM-AE yields the strongest resistance to collective anomalies anomalies (0.903 F1), viable only where continuous GPU provisioning is acceptable.

Future Roadmap & Q&A

Next Steps: Evaluating cross-node federated ensemble voting under strict SWaP constraints and real-world incident labeling.

Requirement: Maintenance of strict chronological separation during lag feature generation to prevent validation leakage.

End of Telemetry Sequence. Questions & Architecture Review.