Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

Applying Unsupervised Machine Learning to Enhance Space Domain Awareness

Vincent Vo, Dylan Treat, Joshua Bagley, Jackson Congrove, Zachary Hruby, and Christopher Gramling

Department of Systems Engineering, United States Military Academy, West Point, New York 10996

Corresponding author's email: christopher.gramling@westpoint.edu

Author Note: The views expressed herein are those of the authors and do not reflect the position of the United States Military Academy, the Department of the Army, or the Department of War. This research was conducted with support from United States Space Command.

Abstract: Space Domain Awareness (SDA) requires the ability to characterize satellite behavior beyond traditional orbital parameter tracking, particularly in the congested and contested geosynchronous (GEO) orbital regime. This study integrates electro-optical (EO) and passive radio frequency (RF) observations from the Unified Data Library (UDL), derived from the Joint Commercial Operations (JCO) cell to construct behavioral features that capture spacecraft maneuver dynamics, attitude, and communications activity. Fused multi-sensor data were conditioned, normalized, and engineered into operationally meaningful indicators including significant changes in visual magnitude, visual magnitude variability, bandwidth variance, frequency variance, and signal-to-noise ratio variance. Satellites were assessed across temporal intervals to evaluate cluster persistence and transitions as indicators of behavioral change and loss of predictability. Results demonstrate that multi-sensor clustering enables scalable behavioral differentiation across mission types and supports operational SDA.

Keywords: Space Domain Awareness, Multi-Sensor Data Fusion, Satellite Behavioral Characterization, Unsupervised Machine Learning, Heterogeneous Multi-Sensor Phenomenology, Visual Magnitude

1. Introduction

There is a growing need for a framework that characterizes satellite behavior over time rather than through isolated events. Many geosynchronous satellites are publicly categorized by broad mission labels, but those labels do not fully capture how satellites behave over time or how they are operationally employed. Increasing competition in the space domain requires the United States to strengthen its ability to monitor and characterize satellite behavior. Pattern of life (PoL) analysis characterizes satellites based on consistent behavior patterns. Traditional SDA relies on maneuver detection, which cannot fully capture stakeholder needs (Kirkpatrick & Shelton, 2024). Unsupervised machine learning has proven effective for pattern-of-life characterization using unlabeled data (Nabi et al., 2024; Pavur & Martinovic, 2021). K-means clustering was selected for this research due to its scalability and interpretability, as supervised methods require labeled data that is often unavailable (Caso, 2025). This project evaluates whether heterogeneous sensor data can characterize satellite behavior over time and identify consistent patterns to support USSPACECOM. Using unclassified GEO data from the Unified Data Library, satellites with complete multi-sensor observations were analyzed. The study also evaluated how time windows and observation thresholds affect coverage. The resulting dataset was clustered using K-means to demonstrate a scalable methodology for identifying satellites that deviate from established behavioral norms.

1.1 Related Work

1.1.1 Maneuver Detection in Geosynchronous Orbit (GEO)

Characterizing satellite behavior in geosynchronous orbit has focused on detecting and classifying orbital maneuvers. Traditional methods rely on orbit propagation and state vector deviation analysis to identify burns associated with station keeping or orbit changes (Roberts, 2023). These approaches detect deviations from predicted trajectories and form the foundation of operational SDA. Unsupervised machine learning has been applied to longitudinal position histories to model Pattern of Life nodes. Roberts et al. used K-means clustering to group drift behaviors and identify transitions such as drift initiation and termination (Roberts, 2023). In these models, behavior is inferred from positional change and maneuver cadence. These approaches remain orbit state driven. Operational changes that do not produce measurable orbital displacement, such as transmission shifts or attitude adjustments, may not be captured. This limitation supports the need for behavioral characterization beyond orbital parameters, which this project's research aims to explore through non-maneuver behavioral analysis.

1.1.2 Phenomenology-Driven Behavioral Characterization in GEO

Recent research has incorporated EO and RF observations into clustering frameworks to characterize satellite behavior. One approach applies k-means clustering using features including delta position, delta velocity, visual magnitude, frequency, bandwidth, time between maneuvers, and time between changes in visual magnitude (El Ouadi et al., 2025). These approaches combine maneuver-derived metrics with photometric and RF signal characteristics to group satellites by operational behavior.

However, several limitations remain. Many existing methods rely on maneuver detection or orbital state estimation features such as delta position and delta velocity. These variables require access to verified maneuver data or high-quality state estimation, which may not be available for all satellites. As a result, maneuver-driven approaches can limit analysis to satellites where such information is observable or accessible. In contrast, observable sensor phenomenology such as EO visual magnitude and RF signal characteristics can be collected for satellites operated by any nation, enabling broader behavioral monitoring.

Additionally, prior work often treats visual magnitude change as a discrete event rather than as a statistical property of brightness variability (Hall, 2023). The present study instead models statistical variability in visual magnitude, bandwidth, frequency, and signal-to-noise ratio across defined time intervals. By removing dependence on maneuver detection and instead modeling statistical variability in observable EO and RF features, this study provides a scalable framework for characterizing satellite behavior across a broader population of space objects.

2. Methodology

The following section describes the methodology used to characterize satellite behavior patterns, using fused data structures from EO and passive RF observations. The workflow of the study begins with data acquisition from the UDL, followed by data fusion and conditioning, feature selection and engineering, unsupervised clustering, and temporal cluster analysis across various time windows. This streamlined approach was designed to characterize baseline satellite behavior and identify deviations from established behavioral patterns over time.

2.1 Data Acquisition, Fusion and Conditioning

Data utilized in this study was acquired from two primary sensor sources available through the United States Space Command Unified Data Library: EO observations and Difference of Arrival (DOA) passive RF measurements. The initial dataset consisted of observations collected during July 2025 across satellites operating in the geosynchronous orbital regime. Two primary sensor sources were used: an EO dataset and a DOA passive RF dataset. These datasets provided complementary observations of satellite brightness behavior and radio frequency transmission activity.

To expand the available observation base, the EO and DOA datasets were extended to include observations from October through December of 2025. For each analysis window, EO observations were fused with DOA detections corresponding to the same satellites and time period. The resulting dataset required conditioning prior to feature engineering due to sensor noise, duplicated detections, and incomplete observations. Preprocessing steps included temporal normalization of observations, removal of duplicate detections, and imputation of missing values where limited EO measurements prevented calculation of normalized visual magnitude features. Missing values were replaced using the mean value across satellites for that feature to preserve dataset completeness. Because only one of five equally weighted features had missing values (~43% of satellites), the impact of imputation was limited.

The fusion process substantially reduced the analyzable satellite population due to the requirement that satellites be observed by both sensing modalities. The initial July dataset contained 1,413 satellites, consisting of 1,365 in the EO dataset and 48 satellites in the DOA dataset. After utilizing cross-sensor availability and integrating the following months' sensor data, the final fused dataset consisted of 48 satellites used for clustering analysis. EO and passive RF data fusion ensured that feature generation reflected operational behavior rather than artifacts introduced by sensor geometry or incomplete data coverage.

2.1.1 Analysis of Time Windows, Observation Threshold, and Coverage

Behavioral characterization requires defining a temporal window over which satellite observations are aggregated. Although the initial dataset discussed previously used July 2025 observations to construct the baseline clustering dataset, the methodology developed in this study is intended to support analysis across multiple observation periods. As a result, selecting an appropriate time window and minimum observation threshold is necessary to ensure sufficient data coverage while preserving sensitivity to behavioral changes. A 90-day dataset from October to December 2025 was used to construct this project's observation threshold graphs, which show the tradeoffs in coverage between different time windows. A smaller time window, for example, 1 week, will have less data but is arguably more useful because a change in behavior can be identified quicker. A larger time window, 1 month, for example, will have more data, therefore making a more accurate characterization, but at a cost of a slower turnaround to the analyst or end user. To visualize this trade-off, a "coverage" score was determined by calculating the percentage of satellites within a dataset that met a given observation threshold for a specified time window. For example, if satellites meet an observation threshold in all months, coverage is 100%; if only one of three months meets the threshold, coverage is 33.3%. Figure 1 shows that this project's threshold effect from 0-10000 observations. The 1-month time window has 80% coverage, where the 1-week and 2-week windows are at 65%. For this study, the 1-month time window was chosen because it maximizes the number of satellites as well as the quality of the characterization and because it has proven to have the best balance between satellite retention, observations, and temporal interpretability.

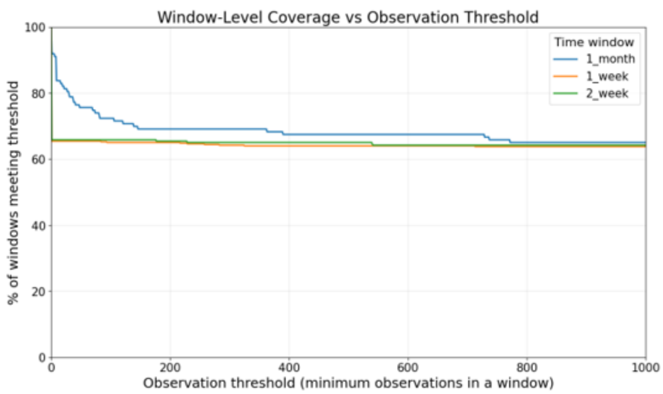

Figure 1. Window-level coverage as a function of observation threshold across three-time windows (1 week, 2 weeks, and 1 month).

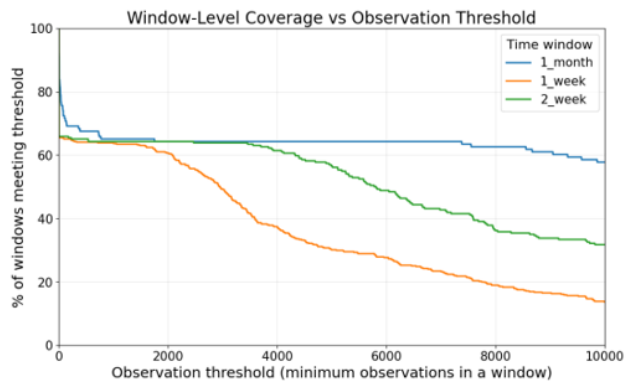

Figure 2. Window-level coverage versus observation threshold for the low-threshold regime (0–1000 observations).

2.2 Feature Engineering and Selection

Feature selection was guided by academic literature in spacecraft photometry, satellite communications, and RF signal analysis, as well as consultation with subject-matter experts from the JCO cell (Beer & Simon, 2021; Leffke et al., 2024; Mariani et al., 2023; Wiersema et al., 2022). The final feature set consisted of five features, two derived from EO observations and three derived from DOA, also known as passive RF sensing:

1. Significant Changes in Visual Magnitude (EO)
2. Visual Magnitude Variability (EO)
3. Bandwidth Variance (RF)
4. Frequency Variance (RF)
5. Signal-to-Noise Ratio (SNR) Variance (RF)

These features were engineered from observational time-series data collected through EO and passive RF sensing architectures. Each feature was selected for its ability to indicate satellite behavior and to capture variability in satellite attitude behavior or communications activity, enabling behavioral characterization across the observed satellite population.

2.2.1 Visual Magnitude Significant Changes and Variability

Visual magnitude variability quantifies changes in how bright a satellite appears when observed by EO ground sensors. Variability in brightness captures how consistently a satellite presents itself optically, and changes in variability are potentially associated with changes in spacecraft attitude, solar panel orientation, or maneuver activity (Mariani et al., 2023; Wiersema et al., 2022). In this study, magnitude variance and change-point detection were used to capture both gradual brightness variability and discrete attitude adjustments across observation windows. Change point detection was performed using pruned exact linear time (PELT).

2.2.2 Bandwidth Variance

Bandwidth variance measures temporal variation in the spectral width of detected satellite transmissions. Changes in bandwidth reflect fluctuations in communications activity and payload utilization. Communications satellites typically maintain stable bandwidth allocations, while sensing satellites often exhibit burst transmission behavior during periodic data downlink events (Li et al., 2024). Measuring variance over time therefore provides a behavioral indicator of transmission dynamics.

ISBN: 97819384962-5-7

2.2.3 Frequency Variance

Frequency variance measures temporal variation in the detected satellite carrier frequency. This is a measure of the consistency in the carrier frequency. Therefore, the variation in the carrier frequency is related to the operational configuration and compensation performed by the satellite (Beer & Simon, 2021; Leffke et al., 2024). In this study, the variance in the carrier frequency was used to detect consistency in the carrier frequency.

2.2.4 Signal-to-Noise Ratio Variance

The signal-to-noise ratio is a measure of the ratio between the signal coming off of a satellite and the noise surrounding it. This is a measure of the consistency in the signal. Therefore, the variation in the signal is related to the operational configuration and compensation performed by the satellite (Beer & Simon, 2021; Leffke et al., 2024). In this study, the variance in the signal-to-noise ratio was used to detect consistency in the signal.

2.3 Unsupervised Clustering via K-Means

Satellites were characterized using K-means clustering, an unsupervised machine learning algorithm that partitions observations into groups based on the similarity of their features (Sinaga & Yang, 2020; Roberts & Linares, 2023). K-means clustering iteratively assigns observations to the nearest cluster centroid and updates centroid locations to minimize within-cluster variance. This process results in a series of clusters that exhibit similar behavioral characteristics. K-means clustering was selected as the unsupervised machine learning algorithm for this study because it provides computational efficiency for large datasets, is easily interpretable, and is compatible with unlabeled data. It is also established as precedent in other academic SDA studies (Roberts & Linares, 2023).

3. Results

3.1 Data Description and Cluster Composition

The optimal number of clusters was determined via an elbow plot, which measures the within-cluster sum of squared distances. The elbow method determined that the optimal number of clusters was 4, at which point diminishing returns in cluster compactness were observed. The satellites were divided into four behavioral clusters based on EO and RF variability features. There were 34 satellites in Cluster 1, 11 satellites in Cluster 2, 1 satellite in Cluster 3, and 2 satellites in Cluster 4. This clustering method was able to identify distinct groups of satellites with differing EO and RF variability features (Figure 3). The final merged data set included 48 satellites that were observed within the July 2025 analysis window. Cluster 1 can be characterized by a higher variability in frequency and bandwidth, which indicates higher volatility and greater activity along the electromagnetic spectrum. Cluster 2 can most aptly be described by a high level of variability in visual magnitude as well as a high number of changes in visual magnitude. This indicates a relatively high level of maneuvers and attitude changes. Two satellites appear as outliers in the principal component projection. These observations indicate that the clustering framework captures both common operational regimes and rare behavioral profiles within the fused dataset. Cluster 3 has only one satellite, which has high RF variability and moderate EO variability. Cluster 4 satellites have frequent discrete brightness changes with low magnitude variability.

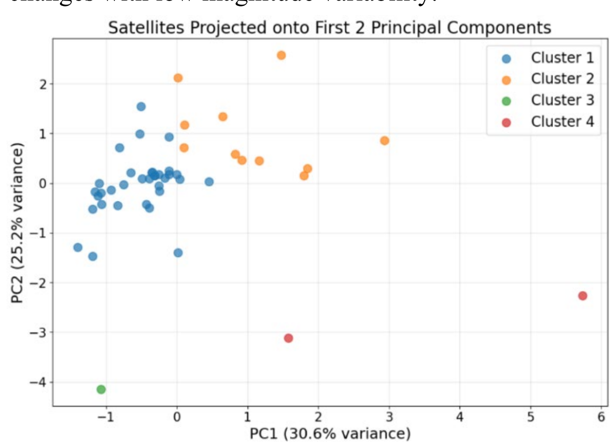

Figure 3. Principal Component Visualization of Clustering Results

ISBN: 97819384962-5-7

3.2 Behavioral Pattern Identification

The pattern of the number of satellites in each anonymized satellite series based on the clusters from the k-means analysis is shown in Figure 3. It can be seen in the figure that most of the satellites are in Cluster 1, followed by a smaller number of satellites in the rest of the clusters. The results show that cluster 1 consisted of seven countries including country A and B, but cluster 2 was comprised of only countries A, B, and C. This suggests that this small subset of satellites in each of these countries are performing a specific mission distinct from most of their other satellites. Within each country, there are several different satellite series, which are groups of satellites with similar payloads and capabilities that perform similar missions. Figure 4 shows the number of satellites from each series in country A and which cluster they fell into. In the available data, country A consisted of seven separate series. All series appeared in cluster 1, but only one of these series appeared in cluster 2. This indicates that country A's series I is behaving in a different manner than all of country A's other satellites, even though they are publicly advertised as performing similar mission sets.

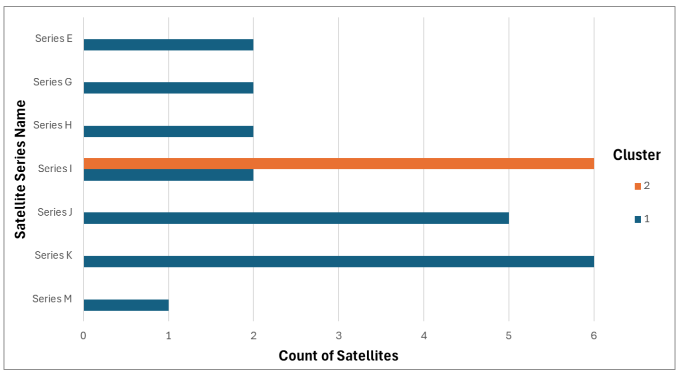

Figure 4. Cluster Results by Satellite Series for Country A

4. Discussion

4.1 Operational Implications for United States Space Command (USSPACECOM)

This project supports Space Domain Awareness (SDA) by providing a scalable framework for behavioral monitoring beyond orbital tracking. Traditional SDA detects positional changes but may miss non-maneuver indicators such as spectrum activity and attitude behavior. The methods used in this project provide several operational advantages, including the ability to identify satellites exhibiting uncharacteristic behavior, scalability to process large satellite populations, and integration of heterogeneous multi-sensor modalities. Incorporating multiple satellite types and data sources into a common behavioral framework improves USSPACECOM's ability to conduct analysis and strengthen SDA. The identification of uncharacteristic behavior is enabled by tracking how satellites move between clusters across different time windows, as repeated changes in cluster assignment may indicate shifts in operational behavior or mission execution. This framework is also scalable, allowing the software to incorporate additional satellite data and potentially operate at higher levels of classification. Multi-sensor fusion strengthens the clustering results by capturing a broader range of behavioral information than any single modality alone. At the same time, preserving satellites of interest during fusion remains important, so future implementations should continue to balance richer fused analysis against the risk of losing relevant satellites due to limited cross-sensor overlap.

4.2 Way Ahead

Future work should expand the fused dataset to improve cross-sensor overlap and statistical robustness. Incorporating additional sensor modalities, including thermal and radar data, could improve behavioral resolution. Refining sensor normalization and feature engineering methods, particularly change point detection and signal stability metrics, may improve interpretability and operational relevance. This framework is intended to support scalable pattern-of-life characterization for satellite behavior assessment and decision-making within SDA environments.

ISBN: 97819384962-5-7

5. Conclusion

This project developed a multi-sensor behavioral characterization framework for satellites using fused unclassified data. Data from the Unified Data Library was processed, engineered into behavioral features, and clustered using K-means to establish baseline patterns of satellite activity. Results show that non-maneuver EO and RF data can characterize satellite behavior over time and identify deviations relevant to Space Domain Awareness. The framework is scalable, reproducible, and adaptable to additional datasets and higher classification levels. While it does not determine adversarial intent, it demonstrates that fused sensor data can reveal meaningful operational changes.

5.1 Acknowledgements

The authors used ChatGPT to assist with editing for grammar, clarity, and conciseness, and incorporated its recommendations into the manuscript.

6. References

Beer, A., & Simon, K. (2021). *Geosynchronous satellite maneuver identification and characterization using passive RF ranging*. https://amostech.com/TechnicalPapers/2021/SSA-SDA/Beer.pdf

Caso, S. (2025). Emerging technologies in military space operations: Current applications and future research for educational and training purposes. *International Journal of Training Research, 23*(2), 140–155. https://doi.org/10.1080/14480220.2024.2431482

Wambui, G. D. (2015). The power of the pruned exact linear time (PELT) test in multiple changepoint detection. *American Journal of Theoretical and Applied Statistics, 4*(6), 581. https://doi.org/10.11648/j.ajtas.20150406.30

El Ouadi, A., Schmedeman, P., & Beskow, D. (2025). Clustering and characterizing geosynchronous satellite behavior. *Procedia Computer Science, 268*, 292–301. https://doi.org/10.1016/j.procs.2025.08.207

Hall, D. T. (2023). Semi-empirical astronomical light pollution evaluation of satellite constellations. *The Journal of the Astronautical Sciences, 69*(6), 1893–1928. https://doi.org/10.1007/s40295-022-00358-4

Kirkpatrick, S., & Shelton, W. (2024). *Space domain awareness as a foundational national security issue*. National Security Space Association. https://nssaspace.org/wp-content/uploads/2024/11/Space-Domain-Awareness-Kirkpatrick-Shelton.pdf

Leffke, Z., Schroeder, K., Phelps, M., & Fletcher, J. (2024). *Passive radio frequency techniques & demonstration for space domain awareness*. https://amostech.com/TechnicalPapers/2024/SDA-Systems-and-Instrumentation/Leffke.pdf

Li, X., Li, Y., Zhao, S., Song, X., & Li, J. (2024). Performance analysis of parallel free-space optical/radio frequency transmissions in satellite–aerial–ground integrated network with power allocation. *Photonics, 11*(12), 1162. https://doi.org/10.3390/photonics11121162

Nabi, I., Farooq, S. Z., Saeed, S., Irtaza, S. A., Shehzad, K., Arif, M., Khan, I., & Ahmad, S. (2024). Leveraging machine learning for the detection of structured interference in global navigation satellite systems. *PeerJ Computer Science, 10*, e2399. https://doi.org/10.7717/peerj-cs.2399

Pavur, J., & Martinovic, I. (2021). On detecting deception in space situational awareness. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security* (pp. 280–291). https://doi.org/10.1145/3433210.3453081

Roberts, T., & Linares, R. (2023). *Geosynchronous satellite behavior classification via unsupervised machine learning*. International Academy of Astronautics Space Traffic Management Conference. https://amostech.com/TechnicalPapers/2021/Machine-Learning-for-SSA-Applications/Roberts.pdf

Roberts, T., Rodriguez, V., Siew, P. M., Solera, H., & Linares, R. (2023). *End-to-end behavioral mode clustering for geosynchronous satellites*. Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS). https://amostech.com/TechnicalPapers/2023/Poster/Roberts.pdf

Roberts, T., Solera, H., & Linares, R. (2023, February 16). *Geosynchronous satellite behavior classification via unsupervised machine learning*. Space Traffic Management Conference. https://www.researchgate.net/publication/368982563_Geosynchronous_Satellite_Behavior_Classification_via_Unsupervised_Machine_Learning

Roberts, T., Solera, H., & Linares, R. (2023, September 21). *Geosynchronous satellite pattern-of-life node detection and classification*. Space Traffic Management Conference. https://www.researchgate.net/publication/368923741_Geosynchronous_Satellite_Pattern-of-Life_Node_Detection_and_Classification

Sinaga, K. P., & Yang, M.-S. (2020). Unsupervised k-means clustering algorithm. *IEEE Access, 8*, 80716–80727. https://doi.org/10.1109/ACCESS.2020.2988796