

Explainable Machine Learning for Space Object Attitude Classification

Marta Guimarães ^{a*}, Maria Almeida ^a, Miguel Santos ^a, Robert Arthur ^a, Chiara Manflettta^a

^a *Neuraspace, Portugal*

*Corresponding Author, marta.guimaraes@neuraspace.com

Abstract

Attitude estimation is a significant challenge in space object tracking. Nevertheless, it also has important applications, including supporting operators facing communication issues and providing a better understanding of object behaviour for trajectory prediction and Space Traffic Management. Traditional attitude estimation methods frequently rely on expert-driven analysis and domain-specific models or require extensive observational data. This work introduces a lightweight and interpretable machine learning approach for attitude determination, leveraging domain-expert-driven feature engineering and decision tree-based models to achieve accurate classification with minimal computational overhead. By designing features that mimic expert decision-making processes, our model is a structured, data-driven alternative to manual analysis, ensuring consistent and reliable classifications – even in scenarios where expert input is unavailable. Our method extracts motion, brightness, and periodicity features from optical observations, effectively distinguishing between tumbling and stable objects. Unlike deep learning-based techniques, which require extensive training and often lack transparency, our approach is interpretable by design and further explainable through *post hoc* analysis methods, reinforcing its reliability for real-world applications. The combination of interpretability and computational efficiency enables rapid experimentation, real-time adaptability, and safe deployment in operational environments. We demonstrate that our method achieves competitive results through extensive validation and ablation studies while maintaining scalability, robustness, and operational feasibility. This work paves the way for efficient and operationally viable attitude estimation, particularly in resource-constrained environments where real-time decision making is crucial.

Keywords: Attitude Characterisation, Optical Observations, Space Traffic Management, Explainable Machine Learning

1. Introduction

Understanding and characterising the attitude behaviour of resident space objects (RSOs) in low-Earth orbit (LEO) has become an increasingly important challenge for the space surveillance and tracking community [1]. In particular, distinguishing tumbling from non-tumbling behaviour can offer valuable insights for object classification, correlation and catalogue build-up and maintenance [2, 3].

Recent work has explored machine learning approaches to address this problem, particularly within the broader context of attitude estimation and characterisation. Two primary sensing modalities have been investigated: radar, using radar cross-section (RCS) measurements, and optical telescopes, using light curves derived from photometric observations. The latter is the focus of our work.

Within the optical passive domain, and especially when applying machine learning, many studies rely on deep architectures to model temporal dependencies in light curves, such as convolutional neural networks [4], or long short-term memory (LSTM) networks [5, 6, 7]. These approaches have shown strong performance but often require dense, high-quality datasets that are not always available in operational settings.

A related example from a different sensing modality is the work of Paulete et al. [8], who proposed a framework for classifying RSOs into various stability modes,

such as Earth-pointing, inertial, or tumbling, based on sequences of RCS observations.

In contrast, our objective is to investigate whether tumbling behaviour can be detected using simpler, more interpretable models trained on sparse data. Rather than modelling fine-grained attitude states, we simplify the problem into a binary classification task: tumbling versus non-tumbling. Inspired by Paulete et al. [8], we derive labels based on object metadata, specifically, the object type and operational status. Although this strategy does not capture the full complexity of tumbling dynamics, it provides a useful proxy that aligns with physical intuition and offers a starting point for scalable attitude characterisation.

To ensure transparency, adaptability, and computational efficiency, we adopt a classical supervised learning pipeline. This includes domain-informed feature engineering [9, 10], and decision-tree-based classifiers, with particular emphasis on model explainability.

Our experiments are based on observational data acquired from Neuraspace's ground-based telescopes, NOWL and SOWL. At the time of the initial experiments, both sensors had been operational for only a few months, resulting in a limited but growing dataset. The data consist primarily of apparent magnitude readings, along with auxiliary pointing information such as azimuth, elevation, and sun phase angle.

2. Methodology

2.1. Approach

This work frames the identification of tumbling objects as a supervised classification problem. Given a set of observational features such as apparent magnitude, sun phase angle, and position of an object in a spherical coordinate system (e.g., azimuth and elevation), the goal is to predict whether an object exhibits tumbling behaviour.

Inspired by the labelling approach proposed by Paulete et al. [8], we construct binary labels based on publicly available metadata. Specifically, objects are labelled as *tumbling* if they are catalogued as debris (DEB) or rocket bodies (R/B), or if they are non-operational payloads (PAY) according to the SATCAT catalogue provided by CelesTrak¹. All remaining objects are labelled as *non-tumbling*, which effectively includes active payloads.

Although this dichotomy simplifies the true diversity of attitude behaviours, it provides a practical and interpretable approach. Non-operational objects and space debris are more likely to exhibit uncontrolled rotational motion due to the absence of active stabilisation mechanisms. In contrast, operational payloads typically maintain some form of attitude control to ensure mission success. While this binary classification label combination can neither fully characterise tumbling dynamics (e.g. tumbling rate or orientation) nor attitude regimes, it provides a meaningful division that serves as a useful proxy for attitude characterisation.

Since the operational status required for this labelling is not available for all objects in LEO, we restrict our analysis to the subset of observations for which this information is known. This decision avoids introducing label noise and ensures that the training data reflects reasonably reliable ground truth. Nonetheless, we acknowledge that the publicly reported operational status may not always be fully up to date or completely accurate, as some objects listed as “operational” may no longer be active, and vice versa. Such discrepancies are difficult to avoid in practice, but given the scope of this initial study, we consider the available metadata a practical proxy for attitude characterisation, with the understanding that refining and validating labels remains an important path for future work.

2.2. Model Choice

As noted, we tackle the challenge of classifying tumbling objects in LEO within a very small data regime. Given the restricted number of observations, our approach emphasises feature engineering and interpretability over entirely data-driven methods.

Beginning with the base features, i.e., right ascension, declination, magnitude, and sun phase angle, we observed that a rule-based approach corresponds closely with expert intuition, as domain specialists frequently

depend on variations in magnitude, the behaviour of phase angles, and the characteristics of motion to evaluate tumbling. This realisation prompted us to implement decision-tree-based models, which inherently capture rule-like patterns and facilitate efficient optimisation.

More concretely, we use XGBoost [11], a state-of-the-art implementation of gradient-boosted decision trees, particularly effective for tabular data. Besides, a key advantage of this approach is its exceptionally quick training time, which facilitates rapid prototyping and experimentation. This flexibility supports ongoing feature refinement, exploration of various resampling techniques for addressing class imbalance, and validation of domain-inspired heuristics.

3. Feature Engineering

Feature engineering plays a crucial role in our study, as we operate within a small data regime where it is vital to maximise the information extracted from each observation. Rather than relying solely on raw observational data (right ascension, declination, magnitude, and sun phase angle), and to leverage the tabular-data approach we have committed to, we have developed a set of features that capture various aspects of brightness variation, periodicity, motion dynamics, and phase angle behaviour of objects in LEO. These features are inspired by physical models of tumbling motion and the heuristic methods traditionally used by domain experts, such as assessing magnitude fluctuations, phase-angle magnitude relationships, and periodic motion patterns.

We categorise the features into five key groups: Magnitude Features, Frequency and Periodicity Features, Temporal Evolution Features, Motion-Based Features, and Phase Angle and Illumination Features. Additionally, we include Date and Seasonality Features to track observational patterns over time.

3.1. Magnitude Features

The magnitude features aim to quantify the variability and distribution of an object’s observed visual magnitude (brightness). Tumbling objects typically exhibit stronger fluctuations in magnitude due to their irregular surfaces and varying orientations relative to the observer and the Sun.

Magnitude Variability Metrics. These features measure how much the object’s magnitude fluctuates over time.

- **mag_range:** Difference between the maximum and minimum observed magnitudes.
- **mag_iqr:** Interquartile range, representing the spread of the middle 50% of the magnitude values.

Magnitude Distribution Shape Metrics. These features analyse the statistical shape of the magnitude dis-

¹celestrak.org/satcat/

tribution to detect asymmetries and extreme fluctuations.

- **mag_skewness:** Measures asymmetry in the magnitude distribution:

$$\frac{1}{N} \sum_{t=1}^N \left(\frac{mag_t - \bar{mag}}{\sigma_{mag}} \right)^3,$$

where mag_t is the observed magnitude at time t , N is the number of observations, $\bar{mag}$ is the mean magnitude, and σ_{mag} is the standard deviation of the magnitude values.

- **mag_kurtosis:** Measures how peaked or heavy-tailed the magnitude distribution is:

$$\frac{1}{N} \sum_{t=1}^N \left(\frac{mag_t - \bar{mag}}{\sigma_{mag}} \right)^4.$$

Scale-Independent Magnitude Variability. These features normalise magnitude fluctuations, making them comparable across objects.

- **mag_cv:** Coefficient of variation, defined as the ratio of the standard deviation to the absolute mean magnitude:

$$\frac{\sigma_{mag}}{|\bar{mag}|}.$$

- **mag_rms:** Root mean square magnitude, which gives greater weight to large fluctuations:

$$\sqrt{\frac{1}{N} \sum_{t=1}^N mag_t^2}.$$

3.2. Frequency & Periodicity Features

Tumbling objects usually have periodic magnitude variations, which can be detected using time-frequency analysis methods. These features help quantify dominant periodicities and power spectral properties.

Lomb-Scargle Periodogram Features. These features quantify periodic behaviour in the frequency domain.

- **ls_dominant_freq:** Frequency corresponding to the highest power peak in the Lomb-Scargle periodogram.
- **ls_power_at_peak:** The power at the dominant frequency, indicating how strong the periodic signal is.
- **ls_power_ratio:** The fraction of total power contained in the dominant period.
- **ls_num_peaks:** The number of significant peaks, indicating multiple periodic components.

Wavelet Transform Features. Unlike Lomb-Scargle, which assumes stationary periodicity, wavelets allow us to analyse how periodicity evolves over time.

- **wavelet_energy:** Total energy in the wavelet transform, representing overall periodicity strength:

$$\sum |W(a,b)|^2,$$

where $W(a,b)$ are the wavelet coefficients at scale a and time b .

- **wavelet_entropy:** Measures how disordered the wavelet power distribution is. Higher values indicate more irregular, chaotic tumbling:

$$-\sum p_i \log p_i,$$

where p_i is the normalised wavelet power at frequency i .

Autocorrelation Features. Autocorrelation measures how similar magnitude values are over time, helping detect repeating patterns.

- **autocorr_lag1:** Correlation between brightness at time t and $t + 1$:

$$\frac{\sum (mag_t - \bar{mag})(mag_{t+1} - \bar{mag})}{\sum (mag_t - \bar{mag})^2}.$$

- **autocorr_lag2:** Correlation between brightness at time t and $t + 2$.

3.3. Temporal Evolution Features

Tumbling objects exhibit changes in magnitude over time, not just in terms of magnitude range but also in how fast and how smoothly magnitude transitions occur. Non-tumbling objects tend to show gradual, predictable changes in magnitude, while tumbling objects may have rapid, erratic, or non-uniform variations. These features capture the rate of change, smoothness, and curvature of the brightness time series.

Magnitude Change Over Time

- **mag_change_max:** Largest observed magnitude change between consecutive observations.
- **mag_change_mean:** Average absolute magnitude change.
- **mag_change_std:** Variability of magnitude change rates.

Magnitude Trend Analysis

- **mag_slope:** The slope of a linear regression fitted to the magnitude time series, indicating whether the object is gradually brightening or dimming:

$$\frac{\sum (t - \bar{t})(mag_t - \bar{mag})}{\sum (t - \bar{t})^2}.$$

- **mag_curvature:** The second derivative of the magnitude time series with respect to time, measuring the rate at which the magnitude change is itself changing (i.e., brightness acceleration):

$$\frac{d^2 mag}{dt^2}.$$

3.4. Motion-Based Features

Beyond magnitude variations, tumbling objects also exhibit irregular motion patterns in their right ascension and declination due to their changing orientation. By analysing the angular velocity, acceleration, and dispersion of motion, we can quantify these variations and distinguish stable objects from tumbling ones.

Angular Speed Features. These features measure how quickly an object's apparent position changes in the sky:

$$\Delta\theta = \sqrt{(\Delta\alpha \cos \delta_{\text{mean}})^2 + (\Delta\delta)^2},$$

$$\delta_{\text{mean}} = \frac{\delta_t + \delta_{t+1}}{2}, \quad \omega = \frac{\Delta\theta}{\Delta t},$$

where $\Delta\alpha$ and $\Delta\delta$ are the changes in right ascension and declination between two observations, δ_{mean} is their average declination, $\Delta\theta$ is the angular displacement, and ω is the angular speed.

- **angular_speed_max:** Maximum recorded angular speed, representing the fastest motion observed.
- **angular_speed_mean:** Average angular speed over the observation period.
- **angular_speed_std:** Standard deviation of angular speed, capturing motion variability (e.g., higher values suggest erratic movement).

Higher-Order Motion Features. These features describe how quickly an object's motion changes, capturing more abrupt shifts in direction or speed.

- **jerk_max:** The maximum rate of change of angular acceleration (i.e., jerk), which reflects how abruptly the motion varies:

$$\frac{\Delta\omega}{\Delta t}.$$

- **motion_dispersion:** The variance of angular speed over time:

$$\frac{1}{N} \sum_{i=1}^N (\omega(t) - \bar{\omega})^2,$$

where $\bar{\omega}$ is the mean angular speed and N is the number of time steps.

Azimuthal Motion Features. These features quantify motion along the azimuth angle, which corresponds to the horizontal direction of the object:

$$\frac{d \text{ azimuth}}{dt}.$$

- **azimuth_speed_max:** Maximum speed in azimuthal direction.
- **azimuth_speed_mean:** Mean azimuthal speed.
- **azimuth_speed_std:** Standard deviation of azimuthal speed.

Elevation Motion Features. These features describe motion along the elevation angle, which corresponds to the vertical position of the object in the sky:

$$\frac{d \text{ elevation}}{dt}.$$

- **elevation_speed_max:** Maximum elevation speed.
- **elevation_speed_mean:** Mean elevation speed.
- **elevation_speed_std:** Standard deviation of elevation speed.

3.5. Phase Angle & Illumination Features

These features analyse how an object's brightness varies with respect to its sun phase angle, which influences how sunlight reflects off its surface and is observed from Earth.

Phase-Magnitude Correlation. Quantifies the linear correlation between observed magnitude and phase angle:

$$\frac{\sum (\theta_t - \bar{\theta})(mag_t - \bar{mag})}{\sqrt{\sum (\theta_t - \bar{\theta})^2 \sum (mag_t - \bar{mag})^2}},$$

where θ_t is the phase angle at time t , mag_t is the observed magnitude, $\bar{\theta}$ is the mean phase angle, and $\bar{mag}$ is the mean magnitude.

Phase Curve Fit Analysis. These features evaluate how well a simple polynomial model fits the phase-magnitude relationship, helping detect deviations from expected reflective behaviour.

- **phase_curve_rmse:** The root mean squared error of the fitted curve:

$$\sqrt{\frac{1}{N} \sum_{i=1}^N (mag_i - \hat{mag}_i)^2},$$

where mag_i is the observed magnitude, $\hat{mag}_i$ is the predicted magnitude from the model, and N is the number of observations.

Table 1. Classification performance of the proposed model over 10 independent trials, reported as mean $\pm$ standard deviation.

Features	F1-Score	Precision	Recall
All	0.641 $\pm$ 0.033	0.643 $\pm$ 0.025	0.642 $\pm$ 0.052
w/o Magnitude Features	0.612 $\pm$ 0.025	0.603 $\pm$ 0.026	0.622 $\pm$ 0.038
w/o Frequency & Periodicity Features	0.617 $\pm$ 0.040	0.569 $\pm$ 0.043	0.676 $\pm$ 0.046
w/o Temporal Evolution	0.618 $\pm$ 0.038	0.613 $\pm$ 0.050	0.628 $\pm$ 0.061
w/o Motion-Based Features	0.552 $\pm$ 0.024	0.562 $\pm$ 0.032	0.545 $\pm$ 0.031
w/o Phase Angle & Illumination	0.613 $\pm$ 0.030	0.606 $\pm$ 0.039	0.624 $\pm$ 0.049
w/o Date and Seasonality	0.575 $\pm$ 0.045	0.543 $\pm$ 0.033	0.613 $\pm$ 0.071

- `phase_curve_variance`: The variance of residuals from the fitted phase-magnitude curve:

$$\frac{1}{N} \sum_{i=1}^N (mag_i - \hat{mag}_i)^2.$$

3.6. Date & Seasonality Features

Finally, we include temporal features. Unlike the other feature groups, this set of features does not directly measure tumbling behaviour but can still provide valuable contextual information about the observational conditions. These features can help account for seasonal variations, observation biases, or long-term trends in the data.

- `obs_month`: Month of observation (1–12).
- `obs_year`: Year of observation.
- `day_of_year`: Day of the year (1–365).
- `time_since_first_obs`: Days since the first recorded observation in the dataset.

4. Model Training and Validation

As outlined in Section 2.2, the classification of tumbling versus non-tumbling objects is performed using XGBoost, a widely used gradient boosting framework known for its performance on structured data. Due to the relatively small dataset and the significant class imbalance, a number of strategies were adopted to enhance the model’s robustness and generalisation.

Handling Class Imbalance with SMOTE. The dataset exhibits a strong class imbalance, with significantly fewer tumbling objects compared to non-tumbling ones. To mitigate this, we apply the Synthetic Minority Over-sampling Technique (SMOTE) [12], which generates new samples for the minority class by interpolating between existing instances. This helps the classifier better learn the decision boundary for the underrepresented class and reduces bias toward the majority class.

Bayesian Optimisation for Hyperparameter Tuning.

XGBoost’s performance is heavily influenced by hyperparameters. To optimise model performance, we employ Bayesian Optimisation [13] to tune key hyperparameters, including maximum tree depth, learning rate, subsample ratio, column sampling ratio, and number of estimators. Unlike grid search, Bayesian Optimisation explores the hyperparameter space more efficiently by using a probabilistic surrogate model to guide the search. Each optimisation run uses 20% of the training set as a validation subset. The F1-score is used as the objective function, as it balances precision and recall (crucial in imbalanced classification tasks). The best combination of hyperparameters is then retained for final model evaluation.

Training and Evaluation. The final model is trained using the full training set with the optimised hyperparameters and the engineered features (see Section 3). Performance is evaluated on a held-out test set using the F1-score, precision, and recall. To ensure robustness, the entire pipeline is executed across multiple random seeds for train/test splits and hyperparameter tuning.

5. Results

For robustness and a fair assessment of the performance, the results were obtained from 10 different trials. Table 1 shows the mean and standard deviation of the F1-score, precision, and recall across all trials when considering different groups of features.

With all the features, the model achieved an average F1-score of 0.641 ± 0.033 , with precision and recall of 0.643 ± 0.025 and 0.642 ± 0.052 , respectively. These values indicate a balanced trade-off between false positives and false negatives, with no significant bias toward one class (as can be noted by the close agreement between precision and recall). The relatively small standard deviations suggest that the model’s performance is consistent across different random initialisations, despite the small dataset size and inherent class imbalance. It is worth noting, however, that a portion of the misclassifications may arise from limitations in the

labelling itself: for example, operational payloads occasionally exhibit tumbling-like behaviour, while certain debris objects may instead be in steady rotation rather than true tumble. Such ambiguities imply that a more carefully curated dataset could lead to improved results, while also highlighting that there may be an upper bound on achievable performance given the current labelling scheme.

The ablation study revealed that removing any feature group leads to a drop in F1-score, confirming that all categories contribute to classification. Motion-based features were the most critical, with their removal causing the largest degradation across all metrics. Interestingly, the highest recall (0.676 ± 0.046) was achieved when frequency- and periodicity-related features were excluded, suggesting that while these features improve precision by filtering false positives, they may also suppress the detection of true tumbling cases. The full set of features, however, achieves the best overall balance between precision and recall.

Although a direct comparison with other works is not possible due to differences in data sources, observation modalities, and labelling schemes, the achieved performance is in line with the range of results reported in related studies such as Paulete et al. [8]. In their work, sequences of RCS measurements processed through deep learning models achieved higher scores when multiple observations per object were available. However, our approach operates in a more challenging setting with sparse optical data and an emphasis on interpretability, factors that make these results a promising first step toward operational attitude characterisation using light curves.

The SHAP [14] summary plot in Figure 1 highlights the most important features driving the classification between tumbling and non-tumbling objects. Each point in the plot corresponds to a SHAP value of an instance and a feature. The features are ordered according to their importance, i.e., the most important one at the top. The importance is defined as the mean absolute value of the SHAP values for each feature. The colourmap represents the feature value in a scale from low to high, i.e., from the minimum feature value to its maximum. For each feature, the jittered points along the y-axis represent the overlapping points to provide a better understanding of the distribution of the values.

Indeed, when analysing Figure 1, extracted when considering all the features, it can be seen that the most important feature is the `time_since_first_obs`, confirming that the temporal context of the observation, specifically, how recently it was acquired, has a strong impact on classification performance. This suggests that the ordering of the samples is not merely incidental but carries meaningful information, with more recent observations potentially reflecting changes in operational status, improved visibility windows or sensor scheduling preferences. Additional temporal metadata, such




Fig. 1. SHAP summary plot of the best performing trial. The importance is defined as the mean absolute value of the SHAP values for each feature. The colourmap represents the feature value in a scale from low to high, i.e., from the minimum feature value to its maximum.

as `obs_month` and `obs_year`, also rank highly, which may capture seasonal visibility patterns, sensor scheduling preferences, or other systematic biases in the dataset rather than intrinsic physical properties of the objects.

The second most important feature, `phase_curve_variance`, quantifies the variance of residuals from a fitted phase-magnitude curve, effectively measuring how much the observed brightness deviates from what would be expected under a simple reflective geometry. High variance indicates that the relationship between brightness and phase angle is poorly explained by a smooth model, suggesting irregular reflective behaviour. Such deviations are physically consistent with tumbling objects, whose changing orientations cause unpredictable changes in observed brightness. Following this, other photometric features such as `mag_change_std` and `mag_rms` also rank prominently, reinforcing the expectation that tumbling behaviour is linked to greater fluctuations in brightness. These features likely capture the compounded effects of rotation, surface heterogeneity, and specular reflections, all of which contribute to more erratic light curves.

Frequency-domain information from `ls_dominant_freq` and `wavelet_energy` contribute further, as expected given that periodic or quasi-periodic signatures in brightness variation are typically associated with tumbling states.

Kinematic features, including `angular_speed_std`, `angular_speed_mean`, and `elevation_speed_std`, also play a relevant role. While their effect on the classification is indirect, they capture physical relationships between observation geometry and apparent magnitude variations, which in turn help distinguish tumbling from stable configurations.

6. Conclusion

We presented a lightweight and interpretable framework for classifying RSOs in LEO as tumbling or non-tumbling from sparse optical observations. By combining domain-expert-driven feature engineering with a decision tree-based model, the method captures physically meaningful motion, brightness, and periodicity features, offering transparency and explainability through SHAP analysis.

Across 10 independent trials, the model achieved an average F1-score of 0.641 ± 0.033 , with precision and recall of 0.643 ± 0.025 and 0.642 ± 0.052 , respectively, indicating balanced and consistent performance despite limited data and class imbalance. While direct comparison with prior studies is not straightforward, the results are competitive with deep learning approaches such as Paulete et al. [8], achieved here under more restrictive data and operational conditions.

This work serves as a starting point, showing that interpretable, feature-based models can achieve promising performance even under sparse observational regimes. The results highlight that such approaches hold potential alongside more complex deep learning solutions, particularly when transparency, explainability, and operational feasibility are priorities.

Acknowledgements

This research was carried out under Project “Artificial Intelligence Fights Space Debris” N° C626449889-0046305 co-funded by the Recovery and Resilience Plan and Next Generation EU Funds (www.recuperarportugal.gov.pt).

References

- [1] Cassien Jobic, Emmanuel Delande, Cristina Perez, Alessandra Dicecco, C Vogt, D Stelmecke, and J Gelhaus. Presentation of the European Union Space Surveillance and Tracking (EU-SST) R&D Plan. In *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2023.
- [2] D. Vallverdú Cabrera, J. Utzmann, and R. Förstner. Integration of attitude characterization in a space debris catalogue using light curves. In *Proceedings of the 8th European Conference on Space Debris (virtual)*, Darmstadt, Germany, 2021. ESA Space Debris Office.
- [3] Laurence Blacketer, Hugh Lewis, and Hodei Urtuxua. Identifying illumination conditions most suitable for attitude detection in light curves of simple geometries. *Advances in Space Research*, 69(3):1578–1587, 2022. ISSN 0273-1177. doi: 10.1016/j.asr.2021.11.010.
- [4] Richard Linares, Roberto Furfaro, and Vishnu Reddy. Space objects classification via light-curve measurements using deep convolutional neural networks. *The Journal of the Astronautical Sciences*, 67(3):1063–1091, 2020.
- [5] Randa Qashoa and Regina Lee. Classification of Low Earth Orbit (LEO) Resident Space Objects’ (RSO) Light Curves Using a Support Vector Machine (SVM) and Long Short-Term Memory (LSTM). *Sensors*, 23(14), 2023. doi: 10.3390/s23146539.
- [6] Ángel Gallego, Carlos Paulete, Marc Torras, Adrián de Andrés, and Alfredo M. Antón. RSO Characterization and Attitude Estimation with Data Fusion and Advanced Data Simulation. In *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2023.
- [7] Gregory P. Badura and Christopher R. Valenta. Physics-guided machine learning for satellite spin property estimation from light curves. *The Journal of the Astronautical Sciences*, 71(5):41, 2024.
- [8] Carlos Paulete, David Cano, Jan Siminski, Cristina Pérez, Diego Escobar, and Jesús Tirado. AIML-RCS: A Machine Learning approach to spacecraft attitude and object identification based on RCS from the S3TSR. In *Proceedings of the 8th European Conference on Space Debris (virtual)*, Darmstadt, Germany, 2021. ESA Space Debris Office.
- [9] Jiří Šilha, Matej Zigo, Tomáš Hrobár, Peter Jevčák, and Martina Verešvárska. Light Curves Application to Space Debris Characterization and Classification. In *Proceedings of the 8th European Conference on Space Debris (virtual)*, Darmstadt, Germany, 2021. ESA Space Debris Office.
- [10] Giorgio Isoletta, Roberto Opromolla, and Giancarmine Fasano. Attitude Motion Classification of Resident Space Objects Using Light Curve Spectral Analysis. *Advances in Space Research*, 75(1): 1077–1095, 2025.
- [11] Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’16*, page 785–794, New York, NY, USA, 2016. Association for Computing Machinery. doi: 10.1145/2939672.2939785.

- [12] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
- [13] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [14] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 2017.