Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Transformer-based Light Curve Inversion for Spacecraft Attitude Classification and Estimation from Photometric Data

Luca Ramponi
University of Arizona
Boris Benedikter
Oklahoma State University
Andrea D'Ambrosio
University of South Florida
Roberto Furfaro
robertof@arizona.edu
University of Arizona
Massimiliano Vasile
University of Strathclyde

Research Article

Keywords:

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11493161/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Transformer-based Light Curve Inversion for Spacecraft Attitude Classification and Estimation from Photometric Data

Luca Ramponi[1], Boris Benedikter[2], Andrea D'Ambrosio[3], Roberto Furfaro[1*], Massimiliano Vasile[4]

[1] Department of Systems and Industrial Engineering, University of Arizona, 1127 E James E Rogers Way,Tucson, 85721,AZ,USA .
[2] School of Mechanical and Aerospace Engineering, Oklahoma State University, 526 N Elgin Ave,Tulsa, 74106,OK,USA .
[3] Department of Mechanical and Aerospace Engineering, University of South Florida, 4202 E Fowler Avenue,Tampa, 33620,FL,USA .
[4] Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street,Glasgow, G1 1XJ,UK .

*Corresponding author(s). E-mail(s): robertof@arizona.edu;

Abstract

As space activities continue to expand, robust and efficient attitude estimation has become central for Space Situational Awareness (SSA). This work introduces a novel transformer-based architecture to tackle the light curve inversion problem for spacecraft attitude and angular velocity determination. By exploiting the self-attention mechanism, the model captures complex temporal dependencies in photometric time series and directly reconstructs spacecraft attitude and angular velocity histories from observed light curves. The approach is trained and tested on a synthetic dataset of light curves generated using a high-fidelity simulator, considering tumbling, spin, nadir-pointing, and Sun-pointing motion regimes, as observed from several ground-based stations. An initial classification stage based on convolutional neural networks (CNNs) distinguishes tumbling and spinning attitude modes from nadir/Sun-pointing regimes, achieving an accuracy of approximately 98%. A subsequent transformer-based regression task reconstructs the full time histories of spacecraft attitude and angular velocity, reaching an average attitude error of less than 1 deg in geodesic distance and an angular velocity magnitude error in the order of 10^{-2} deg/s associated with a direction error of approximately 0.54 deg. In addition, a random forest permutation

importance analysis is carried out to analyze the relative importance of several features, providing insights into which observational and temporal characteristics most strongly influence the model's performance. The results highlight the potential of transformer-based architectures to address the challenges associated with the light curve inversion problem for attitude and angular velocity estimation, marking an advancement in the application of machine learning techniques to complex Space Situational Awareness problems.

Keywords: Attitude Classification, Attitude Estimation, Transformers, Light Curve Inversion, Space Situational Awareness, Machine Learning

1 Introduction

Space Situational Awareness (SSA) is the set of activities that enables the understanding and the characterization of the space environment, including the detection, tracking, monitoring, analysis, and prediction of the behavior of natural and artificial objects in space, thus ensuring the safety and sustainability of space operations [1]. SSA relies on an integrated approach that combines various sensor technologies, ground-based radar systems and optical telescopes, as well as space-based sensors, to provide comprehensive understanding of the space environment.

Each architecture presents unique challenges and trade-offs in terms of coverage, accuracy, cost, and operational complexity. Radar installations deliver precise range and range-rate measurements, independent of solar illumination and weather conditions, but their performance diminishes, and is limited, for high-altitude objects or unfavorable orbital geometries [2]. Phased-array radar systems can also track multiple objects simultaneously, enhancing surveillance capabilities [3]. Nevertheless, the deployment and maintenance of radar systems involve significant costs. On the contrary, optical telescopes are more cost effective and suitable for monitoring objects at higher altitudes due to their lower power requirements, including geostationary orbit (GEO) and the extra-GEO (X-GEO) region, but they are limited to night time observations and clear sky conditions. Among the most effective, space-based sensors provide continuous coverage and are unaffected by atmospheric constraints, but their development, deployment, and maintenance remain technologically and economically demanding. Following these considerations, operational SSA architectures typically rely on a combination of these methodologies to achieve a more general and robust coverage of the space environment.

For the purpose of this work, optical sensors are considered, and especially the two primary types of SSA data they provide, namely astrometric and photometric measurements. Astrometric data refers to the determination of the position of celestial objects by measuring their angular coordinates in time (right ascension and declination). These measurements are crucial for tracking the orbits of satellites and space debris, enabling accurate predictions of their future positions and potential collisions through orbit determination techniques. However, astrometric data alone may not provide sufficient information about the physical characteristics of the observed objects,

such as their shape, size, or attitude which directly affect external perturbations (e.g., atmospheric drag and solar radiation pressure) and consequently their orbital evolution [4]. Photometric data, on the other hand, involves measuring the brightness or intensity of light emitted or reflected by celestial objects over time, resulting in the so-called light curves. Light curves are graphical representations that show how the brightness of an object changes over time, providing insights into its physical properties and rotational state, like shape, material, surface properties and attitude. The combination of astrometric and photometric data enhances the overall understanding of space objects, enabling more accurate orbit determination, attitude estimation, and characterization of their physical properties.

In this work, a transformer-based architecture is proposed to address the light curve inversion problem for the estimation of spacecraft attitude and angular velocity. The objective is to reconstruct the complete time histories of the spacecraft attitude and angular rates using only photometric light curve data and features characterizing it, without requiring any explicit prior knowledge of the spacecraft's physical properties or additional modeling assumptions. This problem falls within the broader framework of light curve inversion, which refers to the process of reconstructing the physical properties and rotational state of a celestial object from its observed brightness variations over time. The concept was first introduced in Ref. [5] to derive details on celestial bodies from opposition phase light curves. Later on, a robust mathematical framework has been developed enabling the approximation of both convex shape and rotational state of asteroids [6]. However, since then, the priority has moved toward developing techniques for artificial objects, such as satellites and space debris, which introduced new challenges due to the complex geometries, heterogeneous construction materials, and surface properties of these objects, which often invalidate the simplifying assumptions previously made for natural bodies. More importantly, artificial objects can have complex attitude dynamics, as they can be subject to a variety of forces and torques that can cause them to tumble or spin in unpredictable ways. Several methods have been proposed to solve the light curve inversion problem for attitude determination of artificial objects, many based on simplifying assumptions regarding the object's shape, material properties, and attitude dynamics. For instance, the space objects have been modeled as simple geometric shapes, such as cylinders for rocket bodies [7], while other methods work by analyzing the light curve as a time signal and extracting features through Fourier decomposition to derive spin period and orientation [8]. Historically, filter-based methods have been widely employed for orbit and attitude estimation [9], most common being the Kalman Filter and its variants. Among these, studies have explored the fusion of astrometric and photometric data within a multiple-model adaptive estimation framework to estimate the most probable shape and inertial orientation of a spacecraft along its trajectory [10]. However, these methods often require accurate initial estimates and can be sensitive to measurement noise and model uncertainties [11]. This can be particularly challenging due to light curves ambiguity, the property for which multiple attitude states can produce similar light curves, especially when including uncertainties in the object's shape and reflective properties, leading to non-unique solutions and possible convergence issues for filter-based methods [12]. As a solution, optimization-based methods have been

proposed, such as Particle Swarm Optimization (PSO), to scan the attitude space and find the best fitting solution. PSOs work by iteratively refining a population of candidate solutions, called particles, based on their fitness with respect to the observed light curve. A first implementation of a PSO of this approach is presented by Clark et al. [13], where the set of possible attitude time histories has been limited heuristically, rather than exploring the whole space of possible solutions, and the method has been tested on slowly rotating objects. A further improvement of the method has been proposed by Burton in Refs. [14, 15] to avoid the need of any initial guess on the attitude state and which is able to solve the ambiguities in the estimation process by exploring the whole space of possible solutions. However, it does require the knowledge of the object's shape, reflective properties, inertia matrix, and the relative positions of the object with respect to the Sun and observer, which are not guaranteed to be known a priori. A recent publication by Gagnon et al. [16] has proposed an approach using a novel multiplicative particle swarm optimization algorithm to optimize the attitude estimation process, which, while giving promising results, still relies on several a priori assumptions on the object's properties. Beyond deterministic optimization and filtering, nonlinear probabilistic estimation frameworks have been introduced to better capture uncertainty and ambiguity in the inversion process. Bayesian inference methods have been applied to jointly estimate attitude, angular velocity, reflective properties, and, in some cases, object shape from photometric data [17]. Markov Chain Monte Carlo techniques, in particular, have been employed to address the strongly nonlinear and ill-posed nature of the problem, enabling shape inversion and attitude estimation within a unified probabilistic framework [18]. While these approaches provide rigorous uncertainty quantification, they are often computationally expensive and depend on sufficiently accurate a priori information. More recently, non-parametric regression techniques have been explored as an alternative probabilistic formulation, such as Gaussian process regression to estimate the attitude of space objects directly from photometric measurements, with unknown surface properties [19]. This approach improves robustness against uncertainties in reflectance parameters, yet the demonstrated application is limited to a single observation window for a GEO object with assumed constant Sun-object-observer geometry. Moreover, the method is not designed to capture the full time history of the attitude state, but rather it estimates the attitude at a single time instant without providing information on the angular velocity of the object, which is instead assumed to be known and used for attitude propagation. Thus, the method does not solve the light curve inversion problem as a complete reconstruction of the attitude dynamics, limiting its applicability to objects with unknown or unpredictable rotational motion.

Machine learning techniques are able to learn complex mappings from data without any prior knowledge, making them well-suited for solving the light curve inversion problem in its most general form. Previous applications demonstrated the ability of deep neural networks to succeed in several SSA related tasks. Specifically, Convolutional Neural Networks (CNNs) have already been shown to work for space object classification from light curve features [20, 21]. Other applications include the synthetic generation of light curves for data augmentation [22], and the characterization of surface composition and attitude motion by combining machine learning

and hyperspectral imaging [23]. Recent developments have also shown the potential of Physics-Informed Neural Networks (PINNs), introduced by Raissi et al. [24], to enforce the physical laws governing the dynamics directly into the learning process. For attitude estimation from light curves, PINN architectures have been proposed to estimate the spin rate of satellites [25], and to use the kinematic equations of attitude motion as a regularization term to constrain the estimation of both attitude and angular velocity time histories [26]. A broader application of large neural networks for space objects behaviour characterization using light curves has been proposed by Groves et al. [27], where a pre-trained perceiver variational autoencoder architecture has been used to build a foundation model for anomaly detection, space object motion prediction, and characterization. Within this framework, transformers, introduced in Ref. [28], have become popular for being used in natural language processing tasks due to their ability to capture long-range dependencies in sequential data, replacing traditional recurrent neural networks (RNNs) and convolutional neural networks in many applications [29, 30]. They rely on a self-attention mechanism that allows the model to capture relationships between different parts of the input sequence, regardless of their distance from each other, and they are not prone to the vanishing gradient problem that can affect RNNs when dealing with long sequences [31]. Other fields of successful application of transformers include computer vision [32, 33], while architecture variants such as the transformer-XL [34] and the gated transformer-XL [35] have been proposed to further improve the ability of transformers to capture longer contextual dependencies and making them more stable and fast to be able to process massive amount of data. Recent research has also explored the application of transformers within the reinforcement learning (RL) framework (meta-RL), aiming to generalize across multiple tasks and environments [36]. Moreover, transformers have also been shown to possess the universal approximation property as operator learning models and successfully applied to forecast solutions of diverse dynamical systems with solutions of finite regularity [37].

Despite their success in other domains, the use of transformer architectures for tackling complex space problems remains relatively limited. Transformers have already been applied to proximity navigation tasks to extract spatial and depth information from simulated images of the International Space Station, demonstrating substantial performance improvements over conventional convolutional neural networks [38]. More recently, transformer-based models have been investigated for spacecraft rendezvous trajectory optimization in supervised learning settings, where they outperformed alternative learning-based approaches in predicting known fuel-optimal trajectories [39]. In parallel, meta-reinforcement learning frameworks incorporating attention-based architectures have been successfully demonstrated to work for space guidance problems, including autonomous fuel-optimal landing scenarios [40]. Beyond guidance and control, transformers have also shown strong performance in space system monitoring tasks. Transformer have been applied to satellite telemetry anomaly detection problems, achieving robust and accurate identification of anomalies, outperforming state of the art methods [41], as well as demonstrated effectiveness in orbit determination tasks [42]. However, to date, transformer-based approaches for astronomical light curve data analysis, particurarly for attitude estimation, are still largely unexplored.

The proposed model exploits the self-attention mechanism to capture complex temporal dependencies and nonlinear relationships between photometric measurements and the corresponding attitude states. A key advantage of the proposed approach is its ability to learn the full sequence to sequence mapping from light curve data to spacecraft attitude dynamics and angular velocity directly from observations, without any other explicit information about the observed satellite. The architecture is trained and validated on a synthetic dataset of light curves generated from several ground-based observatories for both tumbling and spin motion regimes. An initial classification task is performed via CNNs to identify the tumbling and spinning attitude modes from the nadir/Sun-pointing ones, followed by a regression task aimed at reconstructing the full time histories of spacecraft attitude and angular velocity through the proposed transformer model. In addition, a random forest permutation importance analysis is carried out to analyze the relative importance of several features, providing insights into which observational and temporal characteristics most strongly influence the model's performance.

This paper is organized as follows: Section 2 describes the dataset generation process. Details on the methodology are provided in Section 3, including the classification architecture and the transformer architecture employed. Section 4 first reports the outcome of the random forest analysis, then it discusses the results obtained from the proposed approach. Finally, Section 5 draws the main findings of the study and outlines potential future directions.

2 Dataset Model

This section summarizes the generation of the synthetic dataset used to train and validate the proposed transformer architecture for estimating spacecraft attitude and angular velocity time histories from light curves. Photometric measurements have been provided by GMV, generated using their high-fidelity light curve rendering tool, ***Grial***, described in detail in Refs. [43, 44]. The simulator computes the contribution of each illuminated and visible surface element to the reflected irradiance, accounting for its optical properties and self-shadowing between different parts of the object. Rendering is performed using the OpenGL graphics library, allowing the contributions of individual surface pixels to be evaluated through graphics processing unit (GPU) acceleration. The generation process comprises three main computational stages: calculation of the incident solar irradiance; evaluation of surface reflection; and computation of the irradiance received at the observer, followed by its conversion to apparent bolometric magnitude.

First, the solar irradiance (power per unit area) reaching the space object is computed assuming a Sun-object distance of 1 AU and accounting for occultation by the Earth:

$$I_i = \beta \frac{L_\odot}{4\pi r_\odot^2}, \tag{1}$$

where $L_\odot$ is the solar luminosity and $r_\odot$ is the Sun-object distance. The shadowing factor $\beta \in [0, 1]$ is determined from the relative positions of the Sun, Earth, and object, following the formulation described in Ref. [44]. It equals 1 for full illumination and 0 for complete eclipse, with intermediate values representing partial illumination. Here,

I_i denotes the solar irradiance on a plane perpendicular to the incident rays, thus the orientation of each spacecraft surface element is not yet considered. The spacecraft surface is represented as a collection of facets, each characterized by its area, outward unit normal vector, and optical properties. For each illuminated and visible facet, the reflected contribution depends on the unit vectors directed toward the Sun ($\mathbf{l}$) and the observer ($\mathbf{v}$), relative to the facet normal ($\mathbf{n}$). The reflection geometry also involves the half-vector,

$$\mathbf{h} = \frac{\mathbf{l} + \mathbf{v}}{|\mathbf{l} + \mathbf{v}|}, \tag{2}$$

which identifies the microfacet normal orientation required for specular reflection from the Sun toward the observer. The reflection model implemented in ***Grial*** uses a physically motivated bidirectional reflectance distribution function (BRDF), f_r, defined as the ratio of outgoing reflected radiance to incident irradiance for illumination from a single direction. The bidirectional reflectance can be expressed as the sum of a diffuse and a specular component. The specular component $f_{r,s}$, defining the contribution of the mirror-like reflection, is computed using the Cook-Torrance formulation [45], which models each surface as a collection of microfacets while respecting the energy conservation principle:

$$f_{r,s} = \frac{DFG}{4(\mathbf{v} \cdot \mathbf{n})} \tag{3}$$

Here, D is the GGX distribution of microfacet orientations, F is the Fresnel factor and G is the geometric attenuation term. The Fresnel reflectance factor describes the fraction of reflected light from each microfacet, approximated using Schlick's model [46]. The distribution of microfacet orientations is described by the GGX model defined in [47], as it is more accurate for rough materials:

$$D = \frac{n_s^2}{\pi \left[(\mathbf{h} \cdot \mathbf{n})^2 (n_s^2 - 1) + 1\right]^2}, \tag{4}$$

where n_s is the surface roughness coefficient. The geometric attenuation factor G accounts for the self-shadowing and self-masking at the microfacet scale, corresponding to the obstruction of incoming and outgoing light, respectively. Following Ref. [44], it is expressed as

$$G = \min\left\{1, \frac{2(\mathbf{h} \cdot \mathbf{n})(\mathbf{v} \cdot \mathbf{n})}{\mathbf{v} \cdot \mathbf{h}}, \frac{2(\mathbf{h} \cdot \mathbf{n})(\mathbf{l} \cdot \mathbf{n})}{\mathbf{v} \cdot \mathbf{h}}\right\}. \tag{5}$$

The diffuse component is instead described using the Lommel-Seeliger BRDF [48], which accounts for the incidence and viewing geometry to ensure for a more realistic behavior than the Lambertial model. This formulation includes attenuation phenomena due to scattering and absorption when the incident light penetrates the surface. Therefore, the diffuse reflection term is expressed as:

$$f_{r,d} = \frac{\min\{k_d, 1 - F\}(\mathbf{l} \cdot \mathbf{n})}{2\pi \left[(\mathbf{l} \cdot \mathbf{n}) + (\mathbf{v} \cdot \mathbf{n})\right]}, \tag{6}$$

where k_d is the diffuse reflectance coefficient at normal incidence. The factor $\min\{k_d, 1-F\}$ defines the diffuse contribution according to the Fresnel reflectance for a given half-vector.

As a consequence, the irradiance contribution of each visible and illuminated facet γ is computed as:

$$I_{o,\gamma} = (f_{r,s} + f_{r,d})_\gamma \, I_i \, (\mathbf{v} \cdot \mathbf{n})_\gamma \, \frac{A_\gamma}{r^2}, \tag{7}$$

where A_γ is the area of the facet and r is the observer-object distance. The total observed irradiance is obtained as a summation over the surface elements:

$$I_o = \sum_\gamma I_{o,\gamma}, \tag{8}$$

Finally, this measurement is commonly represented in terms of apparent magnitude. Therefore, following the standard conversion:

$$m = -2.5 \log_{10} \left(\frac{I_o}{I_{\text{ref}}} \right) + T_{\text{atm}}, \tag{9}$$

where $I_{\text{ref}} = 2.518 \times 10^{-8}\,\text{W/m}^2$ is the zero point of the apparent bolometric magnitude scale. The term T_{atm} is the atmospheric extinction correction, expressed in magnitudes, computed using the model described in Ref. [49]. It accounts for attenuation by Earth's atmosphere due to Rayleigh scattering, ozone absorption, and diffusion by aerosols.

2.1 Attitude Model

For the purpose of this work, the effect of any external torques is neglected as the observed window is assumed short enough, thus describing a pure kinematic evolution of the attitude state. Thus, considering an uncontrolled space object, a first distinction on its rotational motion can be done by looking at the instantaneous rotation axis. A spinning rotational motion is associated to a rotation axis which is aligned with the object's principal axis of maximum inertia, thus leading to a stable dynamics. In a more general description, non-principal axis rotation involves precession and nutation. Precession is the effect for which the spin axis describes a conical trajectory around the angular momentum direction, fixed in inertial space, while nutation describes periodic variation of the inclination of the rotation axis superimposed on the precession motion. The combination of these effects leads to a more complex rotational state defined as tumbling [50]. In the present work, spinning and tumbling attitude histories are generated using prescribed kinematic laws, without explicitly modeling external torques. The spinning regime is represented by uniform rotation about a specified axis, whereas the tumbling regime is represented by a composition of rotations with specified rotation and precession periods.

The attitude of the space object is represented using quaternions, a four dimensional representation of orientation in three dimensional space that expresses the attitude matrix as a homogeneous quadratic function of the quaternion components, avoiding trigonometric functions and singularities associated with other representations such as Euler angles. However, they require the enforcement of a unit norm

constraint to ensure a valid attitude representation. The formulation employed in this work is last-scalar convention, where the quaternion is defined as $\mathbf{q} = [\mathbf{q}_v, q_s]$, with $\mathbf{q}_v$ being the vector part and q_s the scalar part. Let I denote the inertial frame and B the spacecraft body-fixed frame. The attitude matrix $\mathbf{R}(\mathbf{q})$ mapping inertial frame vector components into body frame components is given by:

$$\mathbf{R}(\mathbf{q}) = \begin{bmatrix} q_1^2 - q_2^2 - q_3^2 + q_4^2 & 2(q_1q_2 + q_3q_4) & 2(q_1q_3 - q_2q_4) \\ 2(q_2q_1 - q_3q_4) & -q_1^2 + q_2^2 - q_3^2 + q_4^2 & 2(q_2q_3 + q_1q_4) \\ 2(q_3q_1 + q_2q_4) & 2(q_3q_2 - q_1q_4) & -q_1^2 - q_2^2 + q_3^2 + q_4^2 \end{bmatrix} \tag{10}$$

Since the available data provided both quaternions and their derivatives, the angular velocity was recovered from the attitude kinematics. Specifcally, let $\boldsymbol{\omega} = \boldsymbol{\omega}_{B/I}^{B}$ be the angular velocity of the spacecraft body frame relative to the inertial frame. Consistently with the scalar-last quaternion convention, the quaternion kinematics is:

$$\dot{\mathbf{q}} = \frac{1}{2}\mathbf{E}(\mathbf{q})\boldsymbol{\omega}, \tag{11}$$

where

$$\mathbf{E}(\mathbf{q}) = \begin{bmatrix} q_4 & -q_3 & q_2 \\ q_3 & q_4 & -q_1 \\ -q_2 & q_1 & q_4 \\ -q_1 & -q_2 & -q_3 \end{bmatrix}. \tag{12}$$

For unit quaternions, $\mathbf{E}(\mathbf{q})^{\mathsf{T}}\mathbf{E}(\mathbf{q}) = \mathbf{I}_3$. Consequently, the angular velocity at each instant is obtained as:

$$\boldsymbol{\omega} = 2\mathbf{E}(\mathbf{q})^{\mathsf{T}}\dot{\mathbf{q}}. \tag{13}$$

3 Methodology

This section presents the methodology developed to address the light curve inversion problem for spacecraft attitude estimation using a transformer-based architecture. The approach is structured in two sequential stages: a classification task followed by a regression task. The first stage performs attitude mode classification, distinguishing among tumbling motion, spinning motions, and a third class that encompasses both nadir-pointing and Sun-pointing configurations. This preliminary classification provides a coarse characterization of the spacecraft's rotational state, supplying contextual information that guides and constrains the subsequent estimation process. It is important to note that only the tumbling and spinning modes are explicitly classified, as the second step is performed only on these attitude configurations. The third class has been included to reduce classification ambiguity, enhance generalization, and improve the robustness of the overall framework. By grouping nadir-pointing and Sun-pointing modes into a dedicated category, the approach prevents light curves with similar characteristics from being incorrectly forwarded to the regression module, which is specifically designed for only the tumbling and spinning motions. The second stage consists of a regression task aimed at reconstructing the complete time histories of the spacecraft attitude and angular velocity directly from the observed light curve data. This task is implemented using the proposed transformer architecture, which is

designed to capture long-range temporal dependencies inherent in photometric time series. The two stages are executed sequentially, with the predicted attitude mode from the classification network incorporated as an additional input feature to the regression model. Figure 1 provides a schematic representation of the overall architecture for a single i-th light curve sample. Here, the input features for both the classification and regression tasks include the light curve magnitude vector $\mathbf{M}_i$ and the relative time vector $\Delta\mathbf{t}_i$ (expressed in seconds from the start of the light curve). Moreover, the second stage takes the absolute time vector in Julian days $\mathbf{t}_{JD_i}$ as an additional input to retain information about the absolute observation epoch. The predicted attitude mode a_i from the classification step is also included as an input feature to the regression model, along with global characteristics extracted from the light curve, such as the mean peak magnitude μ_i and the maximum observed magnitude M_{max_i} (see Section 4.1), as well as the geographical coordinates of the observatory (longitude, latitude, and altitude). The output of the regression model consists of a sequence of predicted quaternions $\mathbf{q}_i$ and angular velocities ω_i time histories. To summarize, the input $\mathbf{x}_{T_i}$ and output $\mathbf{y}_{T_i}$ of the transformer model are:

$$\mathbf{x}_{T_i} = [\mathbf{M}_i, \Delta\mathbf{t}_i, \mathbf{t}_{JD_i}, a_i, lon_i, lat_i, alt_i, \mu_i, M_{max_i}]^T \qquad \mathbf{y}_{T_i} = [\mathbf{q}_i, \omega_i]^T \tag{14}$$

Both architectures are explicitly designed to process padded variable-length sequences, since the light curves in the dataset differ in duration and sampling rate. Accordingly, masking mechanisms are employed to prevent the models from attending to non informative padded elements, ensuring that learning is driven exclusively by valid observational data. Every input feature, except for a_i, is independently standardized (zero mean, unit variance) using the statistics of the dataset to improve the training stability and convergence of the models, as this has been observed to be crucial for the performance of the transformer architecture. The training of both architectures is separately performed using the labeled dataset, therefore the sequential procedure explained above is only implemented during inference, where the predicted attitude mode from the classification step is fed as input to the regression model. This represents a more realistic scenario where the attitude mode of the object is not generally known. The following subsections provide details on the architectures employed for both the classification and regression tasks.

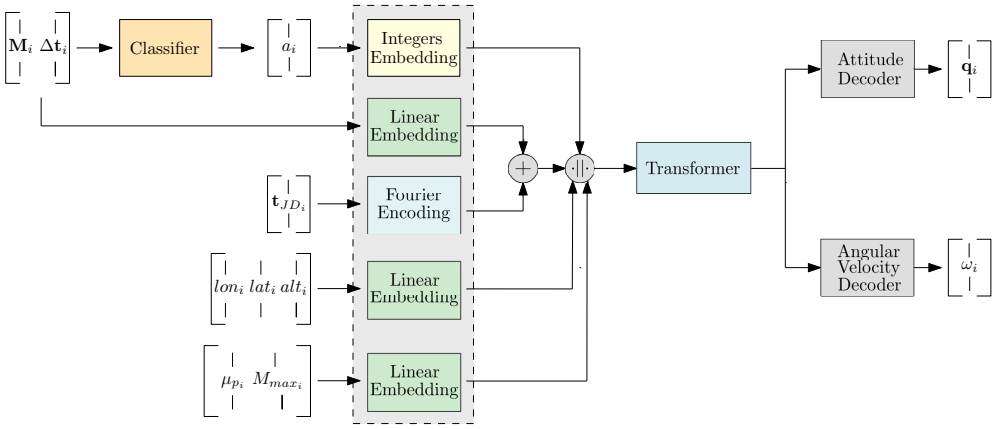

Fig. 1 Schematic representation of the proposed architecture. The architecture consists of a preliminary classification step, followed by a regression step performed by a transformer model

Input Embeddings

Figure 1 shows how every input feature is mapped into an embedding vector of common dimension d_{model}, enabling the transformer to combine photometric measurements, temporal information and global features through self-attention. Different embedding mappings are adopted according to the type of input. The embedding strategy distinguishes between time dependent measurements, represented by one embedding vector per observation, and global features, represented by three additional vectors shared across the observation window.

For the i-th light curve sample, the magnitude and relative time are jointly mapped through a learnable linear projection. Absolute time is represented using a learnable Fourier encoding, which maps a scalar into sinusoidal features with different frequencies:

$$\boldsymbol{\phi}(t) = \left[\sin(\omega_1 t), \ldots, \sin(\omega_{d/2} t), \cos(\omega_1 t), \ldots, \cos(\omega_{d/2} t)\right]^{\mathsf{T}}$$

where $d = d_{\text{model}}$. The Fourier encoding allows to combine lower and higher frequency components, with the latter being more sensitive to small temporal variations. The angular frequencies ω_j are initialized as 2π times uniformly spaced values between 1 and 64, which are also then treated as learnable parameters in the training process. The resulting Fourier encoding is added to the joint projection of magnitude and relative time to form each measurement embedding. This representation combines the elapsed time within the light curve with information about the absolute epoch at which the observations were collected.

Global information about a single light curve sample is included with three additional encoded vectors. Specifically, the observatory coordinates are jointly mapped through one linear projection, while the mean peak magnitude and maximum magnitude are mapped through a separate linear projection layer. However, the attitude

modes are the classification outputs, which are represented by integer indexes. Therefore, they are embedded via the PyTorch module `nn.Embedding`, which allows to map discrete token indices to dense continuous embedding vectors of size d_{model}, which are learned during training. These last three embedding vectors are then concatenated to the rest of the measurement embeddings along the sequence dimension.

3.1 Classifier architecture

As briefly explained above this step is needed as it has been noticed a drastic increase in performance when information about the attitude mode is included in the regression task. However, the attitude mode associated with a light curve cannot be assumed to be known a priori, thus a preliminary classification step is needed to provide this information to the regression model. Convolutional neural networks have already been shown to work well for classification tasks [51, 52], as they are well-suited for processing sequential data and can effectively capture patterns in the data. Therefore, a relatively simple CNN architecture is employed, which takes as input the light curve data and outputs the predicted attitude mode. Specifically, the input is the magnitude of the light curve M_i, which is a one dimensional array of size T, where T is the number of time steps in the light curve and the time duration of the light curve (relative time) Δt_i, where i is the index of the sample in the dataset. The architecture consists of three convolutional layers with ReLU activation functions and dropout regularization layers, followed by a mean pooling over the time dimension and a final fully connected layer. The loss function used for training is the cross-entropy loss implemented in the `cross_entropy` module in Pytorch, which is widely used for classification tasks. Details on the architecture and the training process are provided in Table 1.

Table 1 Classifier architecture details

Hyperparameter	Value
Number of convolutional layers	3
Filters per layer	256, 128, 128
Kernel size	5
Dilation	5
Activation	ReLU
Dropout rate	0.3

3.2 Transformer architecture

Transformers are a type of neural network architecture that relies on self-attention mechanisms to capture long-range dependencies in sequential data. The proposed transformer architecture is designed for sequence to sequence mapping from photometric observations to spacecraft attitude and angular velocity time histories. The goal is to infer the dynamical state evolution of the spacecraft directly from light curve measurements, while accounting for the observational geometry and the prior classification output. The input to the transformer consists of the same features used for the

classification task, namely the magnitude of the light curve and the light curve duration, to which also the absolute time in Julian day as well as the predicted attitude mode from the classification step have been added. In addition, global characteristics extracted from each light curve are included: the mean peak magnitude and the maximum observed magnitude. The peaks are first identified using the `find_peaks` function from the `scipy.signal` library, and the mean peak magnitude is computed as the average magnitude of all detected peaks in the light curve. These features provide statistical information, which help in stabilizing the regression of the dynamical state. Moreover, since the dataset is collected from multiple observatories, the geographical coordinates of each observatory are also fed as input features to provide information about the relative geometry of the observations, which is crucial for accurate attitude estimation. The output of the transformer is a sequence of predicted quaternions and angular velocities time histories of the space object at the same time instants as the light curve measurements. Therefore, following the above description, the architecture takes as input a sequence of feature vectors of dimension $[B \times T \times F]$, where B is the batch size, T is the number of time steps (padded) and F is the number of features (9 in this scenario).

An encoder-only transformer architecture is adopted, as the task involves learning a direct mapping from an input sequence to a corresponding output sequence of the same length, thus an explicit decoding step is not strictly necessary. The encoder is composed of multiple layers, each consisting of a multi-head self-attention mechanism and a position-wise feed-forward network, as described in the "*Attention Is All You Need*" paper [28] and implemented by the `TransformerEncoderLayer` module in Pytorch. The self-attention mechanism allows the model to capture dependencies between different time steps in the input sequence, which is crucial for accurately modeling the complex relationships between the light curve data and the corresponding attitude states.

Multiple encoder layers are then stacked using the `TransformerEncoder` wrapper to construct a deep encoder-only architecture. The padding mask is given as argument to the transformer as `src_key_padding_mask` to ensure that the model correctly handles the variable-length input sequences. In this work, the number of encoder layers, the number of attention heads, the embedding dimension, and the feed-forward network size are treated as hyperparameters that have been tuned to optimize the model's performance on the attitude estimation task. The final transformer design details are provided in Table 2.

Table 2 Transformer architecture details

Hyperparameter	Value
Number of encoder layers	2
Number of attention heads	4
Embedding dimension	128
Feed-forward network size	512

The final encoder output is passed through two separate but identically structured decoder heads, one for the quaternion prediction and one for the angular velocity prediction. Each decoder is a multi-layer perceptron consisting of 2 fully connected layers with a ReLU activation function. It should be noted that the output of the quaternion decoder is normalized to ensure that the predicted quaternions satisfy the unit norm constraint, which is crucial for maintaining a valid attitude representation.

The model is trained using a masked Mean Squared Error (MSE) loss function, which computes the MSE only over the valid time steps of the output sequence, ignoring the padded values:

$$\text{MSE} = \frac{\sum_{i=1}^{N} \sum_{j=1}^{T} m_{i,j} \, (y_{i,j} - \hat{y}_{i,j})^2}{\sum_{i=1}^{N} \sum_{j=1}^{T} m_{i,j}} \tag{15}$$

where N is the number of samples, T is the maximum number of time steps, $y_{i,j}$ is the true value, $\hat{y}_{i,j}$ is the predicted value, and $m_{i,j}$ is the binary mask indicating whether the time step is valid (1) or padded (0). The overall loss function for training is defined as the sum of the MSE losses for the quaternion and angular velocity predictions.

4 Results

The dataset was generated as described in Section 2 for a model of the satellite Sentinel-3A in a low Earth orbit, observed from 28 different observatories. The dataset is composed by 8732 light curves simulated for 4 attitude modes: tumbling (2432), spin (2291), nadir (2321) and Sun pointing (1688). It is split into training and test sets with a 95%-5% split for each attitude mode and each observatory, such that the distribution of attitude modes and observatories is consistent across the different sets. The resulting training set size is equal to 8336 samples, while the test set consists of 396 samples. Among the 396 samples in the test set, 107 are associated with the tumbling attitude mode, 101 with the spin mode, and 188 with the nadir and Sun-pointing modes. Simulations (training and inference) have been carried out using the Python `PyTorch` library on a workstation equipped with Intel(R) Core(TM) i9-14900K, 3200 Mhz, 24 Core(s), 32 Logical Processor(s) and an NVIDIA GeForce RTX 4070 GPU with dedicated memory of 12 GB.

The performance of the classifier is evaluated in terms of recall (R_k) and precision (P_k) for each class, as well as the overall accuracy of the model, defined as:

$$\mathrm{R}_k = \frac{\mathrm{TP}_k}{\mathrm{TP}_k + \mathrm{FN}_k} \qquad \mathrm{P}_k = \frac{\mathrm{TP}_k}{\mathrm{TP}_k + \mathrm{FP}_k} \qquad \mathrm{A} = \frac{\sum_{k=1}^{K} \mathrm{TP}_k}{\sum_{k=1}^{K} (\mathrm{TP}_k + \mathrm{FN}_k)} \tag{16}$$

where TP_k is the number of true positives, FN_k is the number of false negatives, and FP_k is the number of false positives for class k ($k = 0, 1, 2$).

The accuracy on attitude estimation achieved by the transformer regression model is assessed in terms of the angular error in degrees between the predicted and true attitude time histories. This metric, referred to as geodesic distance d_{att}, is computed

through the rotation matrices derived from the predicted and true quaternions:

$$d_{\text{att}} = \cos^{-1}\left(\frac{\text{Tr}(\boldsymbol{R}_{\text{true}}^{T}\hat{\boldsymbol{R}}_{\text{pred}}) - 1}{2}\right) \tag{17}$$

Similarly, the accuracy on the angular velocity estimation is evaluated both as the absolute magnitude error e_{mag_ω} and direction error e_{dir_ω}, defined as the angle in degrees between the predicted and true angular velocity vectors:

$$e_{\text{dir}_\omega} = \cos^{-1}\left(\frac{\boldsymbol{\omega}_{\text{true}} \cdot \hat{\boldsymbol{\omega}}_{\text{pred}}}{\|\boldsymbol{\omega}_{\text{true}}\| \, \|\hat{\boldsymbol{\omega}}_{\text{pred}}\|}\right) \tag{18}$$

4.1 Random Forest Analysis

This section aims at explaining the rationale behind incorporating global characteristic features of the light curves into the transformer input. Due to the complex nature of light curves, influenced by multiple physical and observational factors, it is not straightforward to determine the relative importance of different features in the attitude estimation process. To this end, a random forest algorithm is employed. Random Forest is a machine learning algorithm based on multiple decision trees [53], each looking at different subsets of the data, and is particularly effective for handling complex, large datasets. In this work, a random forest regression model is trained to predict a representative attitude related metric derived from the attitude time history using some predefined features extracted from the light curves. Once trained, the model is used to establish the importance of each feature through a permutation importance analysis. This technique quantifies how much each feature contributes to the model's prediction accuracy by randomly shuffling the values of each feature and measuring the decrease in model performance. Thus, features causing a significant drop in performance are considered to be the most influential. By identifying the most informative factors, it becomes possible to include them into the transformer training process to improve the overall performance. Moreover, it helps in reducing overfitting by encouraging the model to focus on globally informative characteristics rather than on specific patterns in individual data samples.

Let $\mathbf{q}_k$ denote the unit quaternion describing the spacecraft attitude at time step k for the i-th sample. The attitude metric used as the regression target is defined as the mean angular variation over the light curve duration:

$$y_i = \frac{1}{N-1}\sum_{k=1}^{N-1} 2\arccos\left(|\mathbf{q}_k \cdot \mathbf{q}_{k+1}|\right) \tag{19}$$

where N is the number of time steps.

The performance is evaluated in terms of the coefficient of determination R^2, which measures the proportion of variance in the attitude metric that can be explained by

the features extracted from the light curves:

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \tag{20}$$

where y_i is the true value of the attitude metric for the i-th sample, $\hat{y}_i$ is the predicted value from the random forest model, and $\bar{y}$ is the mean of the true values across all samples.

The features considered for the analysis are reported in Table 3, along with their resulting permutation importance mean and standard deviation: the mean and standard deviation of the magnitude of the light curve μ_{M} and σ_{M}, the minimum and maximum observed magnitude $\mathrm{M}_{\mathrm{min}}$ and $\mathrm{M}_{\mathrm{max}}$, the duration of the light curve Δt, the geographical coordinates of the observatory, the number of peaks in the light curve N_{p}, and the mean and standard deviation of the peak magnitudes μ_{p} and σ_{p}. Peaks are detected from the magnitude time series using the `find_peaks` function with default settings. The test is separately conducted for the tumbling and spinning attitude modes to identify potential differences in feature importance between the two regimes.

Table 3 Features considered for the random forest analysis and their resulting permutation importance. Highlighted in gray are the most important features for each attitude mode

Feature	Tumbling	Spin
μ_{M}	0.0665 ± 0.0076	0.0188 ± 0.0051
σ_{M}	0.0517 ± 0.0056	0.0117 ± 0.0034
$\mathrm{M}_{\mathrm{min}}$	0.0438 ± 0.0037	0.144 ± 0.0191
$\mathrm{M}_{\mathrm{max}}$	0.0223 ± 0.0046	0.3963 ± 0.0695
$\Delta\mathrm{M}$	0.2302 ± 0.0178	0.0117 ± 0.0047
Δt	0.3159 ± 0.0294	0.1394 ± 0.0199
lon	0.0063 ± 0.0016	0.0017 ± 0.0015
lat	0.1074 ± 0.0141	0.3294 ± 0.0538
alt	0.0221 ± 0.0046	-0.0033 ± 0.0043
N_{p}	0.0429 ± 0.0043	0.2043 ± 0.0269
μ_{p}	0.5918 ± 0.0306	0.0784 ± 0.0103
σ_{p}	0.0163 ± 0.0034	0.1192 ± 0.0129

The random forest analysis revealed that the mean peak magnitude μ_{p} and the maximum observed magnitude $\mathrm{M}_{\mathrm{max}}$ of the light curve are among the most important features (Figure 2), which is consistent with the physical understanding of the problem, as these features are directly related to the brightness variations of the object, which are influenced by its attitude and rotational state. Other features also show significant importance, but their impact on the attitude estimation process has been observed to be negligible compared to the mean peak magnitude and the maximum observed

magnitude, thus only these two features have been included. Their inclusion in the transformer input allows the model to leverage this global information, improving its ability to learn the complex mapping from light curves to attitude states as will be shown in the results section.

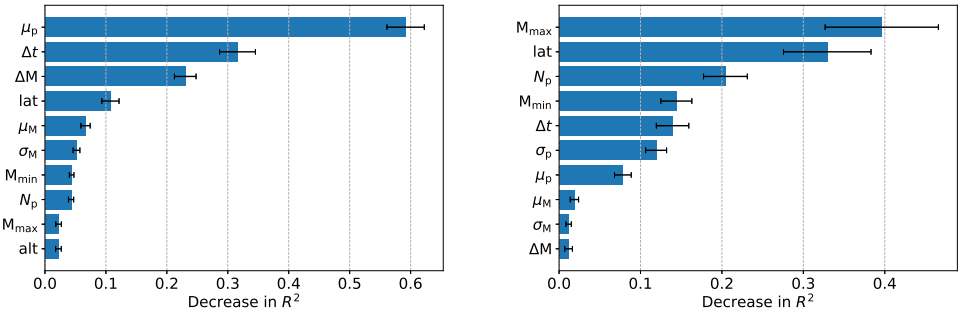

Tumbling mode Spin mode

Fig. 2 Permutation importance of the features for the tumbling and spin attitude modes ordered by importance

4.2 Classification and Regression Results

The classifier architecture is described in Table 1, composed of three convolutional layers with ReLU activation functions and dropout regularization layers, followed by a fully connected layer. The performance in terms of recall and precision for each class, as well as the overall accuracy, are reported in Table 4.

Table 4 Classification results in terms of recall, precision and overall accuracy

Mode	Recall (%)	Precision (%)
Tumbling	100	100
Spin	94.06	96.93
Nadir/Sun-pointing	98.40	96.86
Accuracy (%)	97.73	

The classifier achieves high performance across all classes, with an overall accuracy of 97.73%, demonstrating its effectiveness in distinguishing between the different attitude modes based on the light curve data alone.

As shown in Figure 3, the number of misclassifications is very limited. The few errors arise from confusion between the spin and nadir/Sun-pointing (labeled as N/S-Pointing) classes, which is expected given the similarities in their light curve characteristics compared to the more distinct tumbling mode. Specifically, there are 3 light

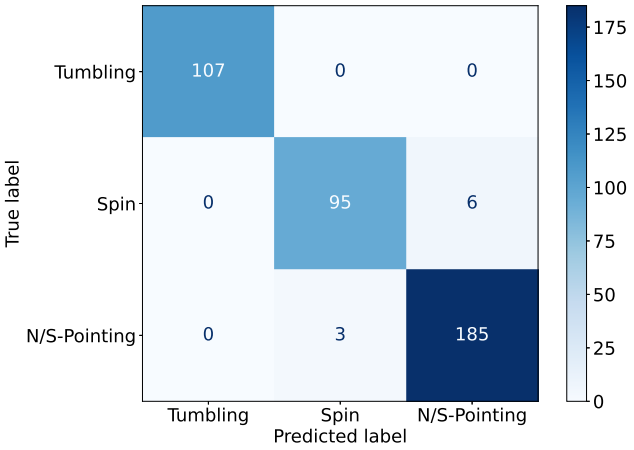

Fig. 3 Confusion Matrix

curves that have been incorrectly classified as spinning, thus leading to non-meaningful regression outputs for those samples.

Therefore, the regression step is performed only on the samples correctly classified as tumbling and spinning, excluding the false positive samples from the evaluation to ensure a fair assessment of the transformer model's ability to estimate attitude and angular velocity time histories, which is the main goal of this work. However, future developments will be focused on improving the robustness of the regression model to misclassifications, as being able to recognize and handle misclassified samples is crucial. A summary of the attitude and angular velocity estimation accuracies is reported in Table 5 in terms of mean and standard deviation for the geodesic distance, as well as the magnitude and direction errors for the angular velocity, while Table 6 provides the same metrics for the tumbling and spinning modes separately.

Table 5 Mean and standard deviation of performance metrics on the test set

	d_{att} (deg)	e_{mag_ω} (deg/s)	e_{dir_ω} (deg)
Mean	0.8987	0.0132	0.5387
Std	0.5647	0.0129	0.5562

The results demonstrate that the proposed transformer architecture accurately estimates attitude and angular velocity time histories from light curve data. The larger error observed in the tumbling mode is attributed to both higher absolute values of

Table 6 Mean and standard deviation of performance metrics per attitude mode

Mode	d_{att} (deg)	e_{mag_ω} (deg/s)	e_{dir_ω} (deg)
Tumbling	1.1485 ± 0.6568	0.0226 ± 0.0112	0.9354 ± 0.4986
Spin	0.6173 ± 0.2064	0.0026 ± 0.0009	0.0919 ± 0.0303

angular velocity and the inherent more complex nature of the motion. On the other hand, the spin mode is characterized by a more regular and predictable motion, which is reflected in the lower estimation errors. Figure 4 illustrates the light curves, attitude, and angular velocity time histories for the worst and median cases across both attitude modes. The best case scenario, associated with a spinning motion, is characterized by low estimation errors, thereby it is not represented as it does not provide meaningful insights into the model's performance given the very high accuracy.

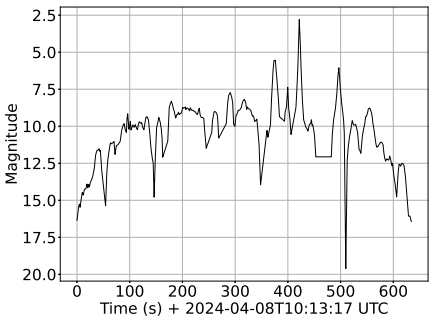

Light Curve Median Case.

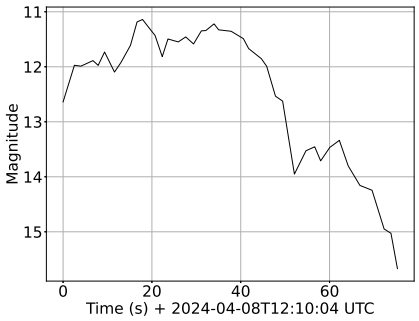

Light Curve Worst Case.

While the previous plots show the model's ability to capture general trends, the estimation accuracy is best quantified through error time histories. Figures 5 and 6 present the absolute attitude and angular velocity errors, respectively, for the best, median, and worst case scenarios. The relative error is also evaluated as the ratio between the absolute error and the corresponding ground truth norm. However, since the attitude is represented using unit quaternions, whose ground truth norm is equal to one, the relative attitude error coincides with the absolute attitude error and is therefore not reported separately. The angular velocity relative error, shown in Figure 7, provides additional insight because its norm varies over time. This highlights that, although the worst case remains the least accurate trajectory, its relative error is approximately one order of magnitude smaller than its corresponding absolute error, reflecting the larger true angular velocity magnitude for that case.

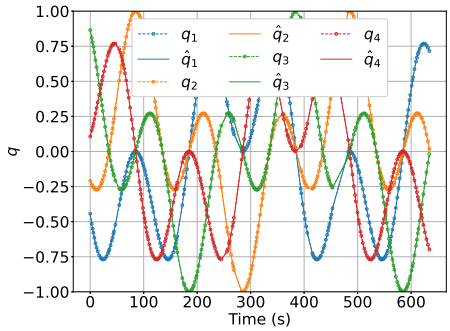

Attitude Median Case

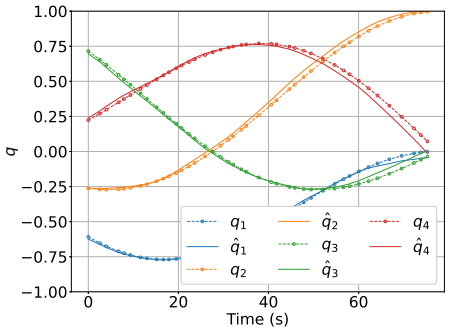

Attitude Worst Case

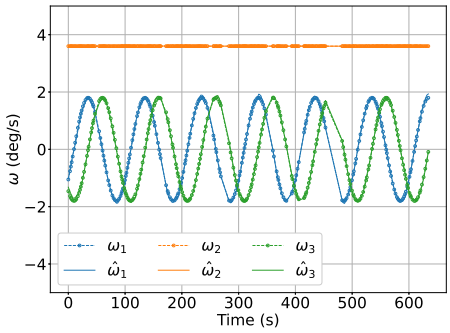

Angular Velocity Median Case

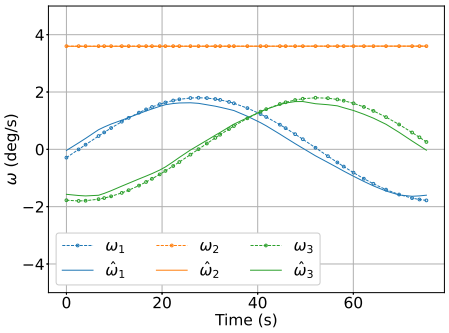

Angular Velocity Worst Case

Fig. 4 Light Curve, attitude and angular velocity time histories for median (left) and worst (right) cases

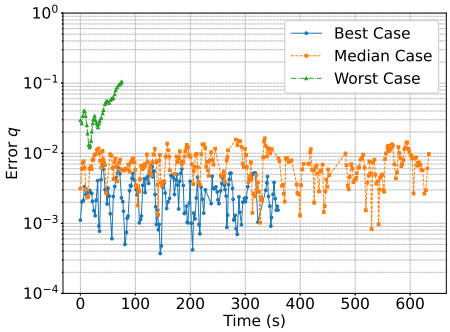

Fig. 5 Attitude absolute error

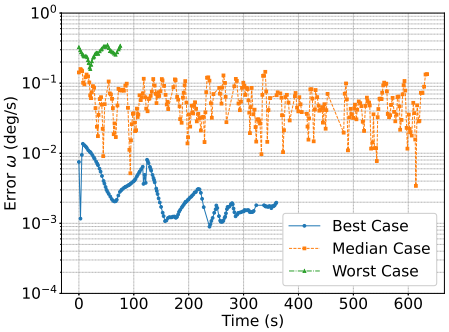

Fig. 6 Angular velocity absolute error

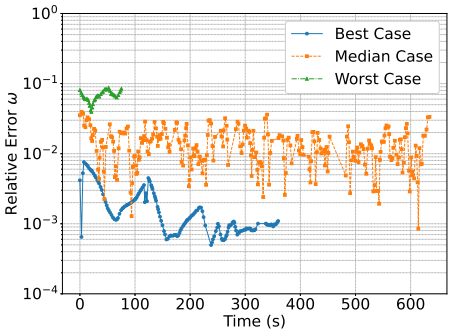

Fig. 7 Angular velocity relative error

The inclusion of short-duration light curves is intended to ensure a more general representation of the dataset, capturing a wide range of observational scenarios. However, it is important to acknowledge that those samples inherently lead to higher estimation errors due to the limited information available. In particular, the worst case scenario corresponds to a light curve with a duration of approximately 75 seconds, which might be insufficient for the model to extract meaningful patterns for accurate attitude estimation. In practical applications, such brief light curves would likely be filtered out or not used for attitude estimation, as they may not provide enough data to support reliable predictions. For instance, the worst case scenario in Figure 4 is associated with an attitude error of 5 degrees, while if a threshold of 100 seconds duration is imposed, the worst case scenario would be of only 2.3 degrees, less than half the error. Additionally, one of the light curves misclassified as a spin mode is less than 100 seconds long, emphasizing the challenges associated with short-duration observation windows. Conversely, the benefits of longer observations are evident in the median case, where a duration of about 600 seconds allows the model to leverage more information, resulting in significantly improved estimation accuracy.

5 Conclusions

This work introduced a novel transformer-based architecture for spacecraft attitude and angular velocity estimation and classification from light curve data. The proposed framework consists of a preliminary classification step to identify the attitude mode, followed by a regression step that reconstructs the complete time histories of the spacecraft's attitude and angular velocity. The architecture is designed to handle variable-length sequences through padding and masking mechanisms, ensuring robust performance across diverse observational scenarios. Moreover, no initial guess of either the attitude, angular velocity or any other satellite characteristic is required as an explicit input to the model, making it suitable for a wide range of space objects and observational conditions. A random forest analysis was conducted to identify the most informative features from the light curves, leading to the inclusion of global characteristics such as mean peak magnitude and maximum observed magnitude in the

transformer input. The results demonstrate the effectiveness of the proposed approach, achieving a classification accuracy of 97.73%. For correctly classified tumbling and spinning samples, the transformer provides accurate attitude and angular velocity estimates, with mean errors of about 0.9 degrees in attitude, 0.5 degrees in angular velocity direction and 0.01 deg/s in angular velocity magnitude. The analysis on the error over time revealed the influence of light curve duration on the estimation accuracy, with shorter observation windows generally associated with larger errors, consistent with the reduced information content available to the model. Although the current framework has produced encouraging results, additional work is required to confirm its validity and robustness. Future improvements will involve incorporating observational noise into the training dataset, exploring a wider range of satellite geometries, and developing a more robust regression model capable of managing misclassifications arising from the initial classification stage. Moreover, the model's ability to generalize will be evaluated using real light curve data, offering deeper insight into its practical performance and possible limitations under real world conditions.

Acknowledgements. This work was supported in part by the UK Space Agency through the International Bilateral Fund, grant UKSAG24-081. The authors gratefully acknowledge GMV (Spain) for providing the light curve data as part of the UK-IBF project.

Statements and Declarations

Funding

This work was supported in part by the UK Space Agency through the International Bilateral Fund, grant UKSAG24-081.

Competing interests

The authors declare no competing interests relevant to this work.

Data availability

The synthetic light curve data used in this study were provided by GMV within the UK-IBF project and are not publicly available.

References

[1] Gasparini, G., Miranda, V.: In: Rathgeber, W., Schrogl, K.-U., Williamson, R.A. (eds.) Space situational awareness: an overview, pp. 73–87. Springer, Vienna (2010). https://doi.org/10.1007/978-3-211-99653-9_7 . https://doi.org/10.1007/978-3-211-99653-9_7

[2] Muntoni, G., Montisci, G., Pisanu, T., Andronico, P., Valente, G.: Crowded space: a review on radar measurements for space debris monitoring and tracking. Applied Sciences **11**(4), 1364 (2021)

[3] Cataldo, D., Gentile, L., Ghio, S., Giusti, E., Tomei, S., Martorella, M.: Multi-bistatic radar for space surveillance and tracking. IEEE Aerospace and Electronic Systems Magazine **35**(8), 14–30 (2020) https://doi.org/10.1109/MAES.2020.2978955

[4] Frueh, C.: AAE 590: Space Traffic Management, Lecture Notes. Purdue University, School of Aeronautics and Astronautics (2023)

[5] Russell, H.N.: On the light variations of asteroids and satellites. The Astrophysical Journal **24**, 1–18 (1906) https://doi.org/10.1086/141361

[6] Kaasalainen, M., Torppa, J., Muinonen, K.: Optimization methods for asteroid lightcurve inversion: Ii. the complete inverse problem. Icarus **153**(1), 37–51 (2001) https://doi.org/10.1006/icar.2001.6674

[7] Hrobar, T., Silha, J., Zigo, M., Jevcak, P., Patole, P., Schildknecht, T.: Attitude determination of cylindrical rocket bodies by using simultaneous bistatic photometric measurements. In: Ryan, S. (ed.) Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, p. 117 (2023)

[8] Benson, C.J., Scheeres, D.J., Ryan, W.H., Ryan, E.V.: Cyclic complex spin rate evolution of defunct geo satellites. In: Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference (2018)

[9] Crassidis, J.L., Junkins, J.L.: Optimal Estimation of Dynamic Systems, 2nd edn. Chapman and Hall/CRC, Boca Raton, FL (2011). https://doi.org/10.1201/b11154

[10] Linares, R., Crassidis, J., Jah, M., Kim, H.: Astrometric and photometric data fusion for resident space object orbit, attitude, and shape determination via multiple-model adaptive estimation. In: AIAA Guidance, Navigation, and Control Conference (2010). https://doi.org/10.2514/6.2010-8341

[11] Wetterer, C.J., Jah, M.: Attitude determination from light curves. Journal of Guidance, Control, and Dynamics **32**(5), 1648–1651 (2009) https://doi.org/10.2514/1.44254

[12] Holzinger, M.J., Alfriend, K.T., Wetterer, C.J., Luu, K.K., Sabol, C., Hamada, K.: Photometric attitude estimation for agile space objects with shape uncertainty. Journal of Guidance, Control, and Dynamics **37**(3), 921–932 (2014) https://doi.org/10.2514/1.58002 https://doi.org/10.2514/1.58002

[13] Clark, R., Dave, S., Wawrow, J., Lee, R.: Performance of parameterization algorithms for resident space object (rso) attitude estimates. In: Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) (2020). https://api.semanticscholar.org/CorpusID:245959950

[14] Burton, A., Robinson, L., Frueh, C.: Light curve attitude estimation using particle swarm optimizers. Advances in Space Research **74**(11), 5619–5638 (2024) https://doi.org/10.1016/j.asr.2024.09.008

[15] Burton, A.: Attitude estimation using light curves (2024) https://doi.org/10.25394/PGS.26383684.v1

[16] Gagnon, S., Crassidis, J.L.: Particle swarm optimization on the space of quaternions with applications to attitude estimation from light curves. The Journal of the Astronautical Sciences **72**(2), 8 (2025) https://doi.org/10.1007/s40295-025-00486-7

[17] Campbell, T., Furfaro, R., Reddy, V., Battle, A., Birtwhistle, P., Linder, T., Tucker, S., Pearson, N.: Bayesian approach to light curve inversion of 2020 so. The Journal of the Astronautical Sciences **69**(1), 95–119 (2022)

[18] Linares, R., Crassidis, J.L.: Space-object shape inversion via adaptive hamiltonian markov chain monte carlo. Journal of Guidance, Control, and Dynamics **41**(1), 47–58 (2018) https://doi.org/10.2514/1.G002296 https://doi.org/10.2514/1.G002296

[19] Hara, R., Yoshimura, Y., Hanada, T.: Attitude estimation from photometric data using gaussian process regression. Journal of Space Safety Engineering **12**(1), 227–238 (2025) https://doi.org/10.1016/j.jsse.2025.02.008

[20] Furfaro, R., Linares, R., Reddy, V.: Space objects classification via light-curve measurements: deep convolutional neural networks and model-based transfer learning. In: AMOS Technologies Conference, Maui Economic Development Board, pp. 1–17 (2018)

[21] Furfaro, R., Linares, R., Reddy, V.: Shape identification of space objects via light curve inversion using deep learning models. In: AMOS Technologies Conference, Maui Economic Development Board, Kihei, Maui, HI (2019)

[22] D'Ambrosio, A., Scorsoglio, A., Battle, A., Furfaro, R., Reddy, V.: Machine learning-based light curves brightness prediction for space objects in the geostationary belt. In: AIAA SCITECH 2024 Forum, p. 1672 (2024)

[23] Marto, S.G., Vasile, M., Campbell, A., Murray, P., Marshall, S., Savitski, V.: Hyperspectral lightcurve inversion for attitude determination. In: AIAA SCITECH 2024 Forum, p. 0198 (2024)

[24] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics **378**, 686–707 (2019)

[25] Badura, G.P., Valenta, C.R.: Physics-guided machine learning for satellite spin property estimation from light curves. In: Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, p. 71 (2023)

[26] Benedikter, B., D'Ambrosio, A., Scorsoglio, A., Furfaro, R., Vasile, M.: Physics-informed machine learning for attitude estimation from light curves. In: AAS/AIAA Space Flight Mechanics Meeting (2025)

[27] Groves, I., Campbell, A., Fernandes, J., Rodríguez, D.R., Murray, P., Vasile, M., Nockles, V.: A Self-Supervised Framework for Space Object Behaviour Characterisation (2025). https://arxiv.org/abs/2504.06176

[28] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. NIPS'17, pp. 6000–6010. Curran Associates Inc., Red Hook, NY, USA (2017)

[29] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., Rush, A.: Transformers: State-of-the-art natural language processing. In: Liu, Q., Schlangen, D. (eds.) Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 38–45. Association for Computational Linguistics, Online (2020). https://doi.org/10.18653/v1/2020.emnlp-demos.6 . https://aclanthology.org/2020.emnlp-demos.6/

[30] Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR **abs/1810.04805** (2018) 1810.04805

[31] Luong, M., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation. CoRR **abs/1508.04025** (2015) 1508.04025

[32] Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. CoRR **abs/2012.00364** (2020) 2012.00364

[33] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. CoRR **abs/2010.11929** (2020) 2010.11929

[34] Dai, Z., Yang, Z., Yang, Y., Carbonell, J.G., Le, Q.V., Salakhutdinov, R.: Transformer-xl: Attentive language models beyond a fixed-length context. CoRR **abs/1901.02860** (2019) 1901.02860

[35] Parisotto, E., Song, F., Rae, J., Pascanu, R., Gulcehre, C., Jayakumar, S., Jaderberg, M., Kaufman, R.L., Clark, A., Noury, S., Botvinick, M., Heess, N., Hadsell, R.: Stabilizing transformers for reinforcement learning. In: III, H.D., Singh, A. (eds.) Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 119, pp. 7487–7498. PMLR, Virtual (2020). https://proceedings.mlr.press/v119/parisotto20a.html

[36] Melo, L.C.: Transformers are meta-reinforcement learners. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) Proceedings of the 39th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 15340–15359. PMLR, Baltimore, Maryland, USA (2022). https://proceedings.mlr.press/v162/melo22a.html

[37] Shih, B., Peyvan, A., Zhang, Z., Karniadakis, G.E.: Transformers as neural operators for solutions of differential equations with finite regularity. Computer Methods in Applied Mechanics and Engineering **434**, 117560 (2025) https://doi.org/10.1016/j.cma.2024.117560

[38] Ghilardi, L., Scorsoglio, A., Furfaro, R.: Iss monocular depth estimation via vision transformer. In: Ieracitano, C., Mammone, N., Di Clemente, M., Mahmud, M., Furfaro, R., Morabito, F.C. (eds.) The Use of Artificial Intelligence for Space Applications, pp. 167–181. Springer, Cham (2023)

[39] Guffanti, T., Gammelli, D., D'Amico, S., Pavone, M.: Transformers for Trajectory Optimization with Application to Spacecraft Rendezvous (2024). https://arxiv.org/abs/2310.13831

[40] Federici, L., Furfaro, R.: Meta-reinforcement learning with transformer networks for space guidance applications. In: AIAA SCITECH 2024 Forum (2024). https://doi.org/10.2514/6.2024-2061 . https://arc.aiaa.org/doi/abs/10.2514/6.2024-2061

[41] Fejjari, A., Delavault, A., Camilleri, R., Valentino, G.: Transformer-based anomaly detection for satellite telemetry data. Acta Astronautica **238**, 739–745 (2026) https://doi.org/10.1016/j.actaastro.2025.09.035

[42] Ré, N.P., Popplewell, M., Caudill, M., Sullivan, T., Hanf, T., Tatman, B., Parmar, K., Presser, T., Chikine, S., Grant, M., Poulson, R.: Transformers for orbit determination anomaly detection and classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pp. 6819–6827 (2024)

[43] Rubio, J., Andrés, A., Paulete, C., Gallego, Á., Lozano, J.M., Escobar, D., Agueda, A.: A high-performance and robust light curve inversion method for attitude monitoring of three-axis stabilised satellites. In: The Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, p. 138 (2025)

[44] Rubio, J., de Andrés, A., Paulete, C., Gallego, Á., Escobar, D.: Attitude estimation of uncontrolled space objects: A bayesian-informed swarm intelligence approach. Advances in Space Research **77**(3), 3791–3814 (2026) https://doi.org/10.1016/j.asr.2025.11.112

[45] Cook, R.L., Torrance, K.E.: A reflectance model for computer graphics. ACM Trans. Graph. **1**(1), 7–24 (1982) https://doi.org/10.1145/357290.357293

[46] Schlick, C.: An inexpensive brdf model for physically-based rendering. Computer Graphics Forum **13**(3), 233–246 (1994) https://doi.org/10.1111/1467-8659.1330233 https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.1330233

[47] Walter, B., Marschner, S.R., Li, H., Torrance, K.E.: Microfacet models for refraction through rough surfaces. In: Proceedings of the 18th Eurographics Conference on Rendering Techniques. EGSR'07, pp. 195–206. Eurographics Association, Goslar, DEU (2007)

[48] Fairbairn, M.B.: Planetary Photometry: The Lommel-Seeliger Law. Journal of the Royal Astronomical Society of Canada **99**(3), 92 (2005)

[49] Buil, C.: Le calcul de la transmission spectrale atmosphérique: application à la surveillance de l'activité Be d'étoiles faibles. http://www.astrosurf.com/aras/extinction/calcul.htm (2009)

[50] Schaub, H., Junkins, J.L.: Analytical Mechanics of Space Systems. AIAA education series. American Institute of Aeronautics and Astronautics, ??? (2009). https://books.google.com/books?id=3FUyQAAACAAJ

[51] Yao, G., Lei, T., Zhong, J.: A review of convolutional-neural-network-based action recognition. Pattern Recognition Letters **118**, 14–22 (2019) https://doi.org/10.1016/j.patrec.2018.05.018 . Cooperative and Social Robots: Understanding Human Activities and Intentions

[52] Dhillon, A., Verma, G.K.: Convolutional neural network: a review of models, methodologies and applications to object detection. Progress in Artificial Intelligence **9**(2), 85–112 (2020) https://doi.org/10.1007/s13748-019-00203-0

[53] Breiman, L.: Random forests. Machine Learning **45**(1), 5–32 (2001) https://doi.org/10.1023/A:1010933404324