RSO Characterization and Attitude Estimation with Data Fusion and Advanced Data Simulation

Ángel Gallego, Carlos Paulete, Marc Torras, Adrián de Andrés and Alfredo M. Antón
GMV

ABSTRACT

Space Surveillance and Tracking (SST) routine operations often involve orbit determination for large catalogues of Resident Space Objects (RSOs), combining multiple data sources to yield more accurate results. However, there is a growing interest in using the same and other sources of data to perform an object characterization that goes beyond orbit determination, providing valuable insights for commercial, civil, and military SST and Space Domain Awareness (SDA) applications.

This paper proposes a methodology for characterizing RSOs by combining data from typical SST and SDA sensors along with information from other external sources. It comprises several algorithms chained in an incremental approach. First, objects are classified according to their type of stabilization, using Machine Learning techniques based on telescope light-curves and Radar Cross Section (RCS) data. Subsequently, for rotating objects, the apparent rotation period is determined by means of Lomb-Scargle periodograms and Epoch Folding, also based on light-curves and RCS. From the same data, an additional method estimates the size of objects when it is unknown. Finally, previous results combined with external data are used within the attitude estimation process, which determines the pointing law for the 3-axis stabilized objects or the rotation axis and rate for objects in rotation. This algorithm processes light-curves based on an iterative Least Squares Method (LSM) filter, which minimizes residuals between real measurements and simulated measurements. It is at this point that a light-curve simulator is required. The one used within this methodology has been developed based on OpenGL and is also presented in this paper.

Although further enhancements can be explored, the results obtained so far, as discussed here, suggest that the methodology has the potential to be highly valuable for a range of future applications.

1. INTRODUCTION

In the last decades, the advances in the various fields of Space Surveillance and Tracking (SST) have been impressive. Among many other aspects, Resident Space Object (RSO) cataloguing methodologies have advanced with novel algorithms and models for data correlation, orbit determination (OD) and propagation, uncertainty estimation and manoeuvre detection. In turn, traditional SST services have also been enhanced with new algorithms and improvements coming from cataloguing (e.g., huge increase of accuracy in risk and collision probability calculations by new methods combined with a strong knowledge of the orbit and its uncertainty). In parallel, sensor networks, both commercial and institutional, have evolved in different types of technologies, in number of sensors, in performance, in coverage, and have increased the variety, types and accuracy of data they can provide.
As a consequence of these improvements in SST, the last few years have seen an increase of thousands of catalogued objects, with a decrease of the lower size limit for their observability, higher accuracies and lower orbital uncertainties. In addition, private and public initiatives and investments in SST systems have multiplied for developing sensors, catalogues and services, which provide their data and help to ensure the safety of space operations.

However, there are areas related to SST, even in the civil environment, where knowledge of the orbit of objects is not enough. For example, knowing the size or shape of an object can be critical for the assessment of a collision or re-entry risk. There are also new services, such as those related to space debris removal, where other aspects, like the attitude of the object to be removed, are a determining factor for the success of the mission.
On the other hand, space is also consolidating as a strategic environment, where it is necessary to have an additional perspective of knowledge of the objects in orbit, for a comprehensive understanding of their behaviour, intentions, and potential threats. This need has given rise to a crucial aspect of space security, the Space Domain Awareness (SDA), which extends beyond traditional SST to encompass a more holistic perspective.

These latter aspects, both in the SDA and SST environment, demonstrate that more is needed than just precise knowledge of the orbit of objects: their characterisation. This intricate and broad concept implies identifying or estimating a varied set of physical and spectral characteristics of RSOs, such as their size, shape, equipment, materials, rotation, attitude, electronic signature... This holistic characterization enables a more in-depth understanding of the objects' missions, capabilities, and intentions, which is of immediate interest in military applications. But not only there, as mentioned previously, as the attitude state of an object is vital for some civil activities, such as active debris removal, and for commercial operations, for example to provide support in contingency cases (e.g. with Safe modes activated) or to collaborate in phases such as Launch and Early Orbit Phase (LEOP) or End-of-Life (EOL) operations.

And to be useful in the SST or SDA domains, it must be done on the basis of the data provided by the sensors used in these fields. More specifically, based on processing the measurements that usually complement the astrometry they provide, such as Radar Cross-Section (RCS) in the case of radars, light-curves for optical sensors, and laser light-curves for laser sensors. But not only this data can be useful, as many of the sensors often used in SST and SDA, especially the more advanced ones, are able to directly provide object characterisation (e.g. size and shape estimation, rotation and attitude information). This being the case, one might think that the only thing to do is to use these sensor data to obtain the above-mentioned characterisation. The problem comes when such data are not available for an object. This can occur because the sensors used are not able to provide such data directly, or because there is a lack of adequate accuracy, coverage or passes of the objects with these sensors. And there could also be problems of lack of accessibility or realisability of such data, which is very common in cases of military sensor networks.
It is also important to note that there are public catalogues containing object characterisation information. This is an extremely useful information for a wide variety of cases, even more considering its open availability. However, these types of external catalogues also have limitations, such as the lack of information for some objects, or the total lack of information for non-public objects (precisely those of interest from a military perspective).
In addition, it must also be considered that these different sources sometimes provide different information, and as the procedures used to derive these data and the associated accuracy are not known, it is difficult to determine which is the most reliable source in each case.

This makes it necessary to look for alternative characterisation methodologies that cover the widest possible data exploitation (e.g. light-curves [1] or RCS [2]), so as not to rely solely on the direct characterisation provided by the sensors or external catalogues. In this regard, Machine Learning (ML) methodologies are often employed to classify objects according to their stabilisation, shape or even directly derive rotation or attitude parameters [3][4]. Although these methods are powerful, the general problem they present is the need to have a large data set available to train the models correctly. Although data can be simulated, they must be representative and of a similar characteristic to the data that will be used in their real application, which means that in many cases simulated data are not suitable for this purpose. Even when enough real data is available for model training, the problem arises that it is often not possible to use the model to apply it to data from other sensors. Therefore, although these methods are a good starting point and useful in sometimes, they are not fully applicable and accurate enough in all general cases.
On the other hand, in the particular case of attitude estimation methods, there are algorithms using iterative filters [9][10] to minimise residuals, which implies the need to have a simulator of measurements of the same type as those being processed, which can be a very complex task. In addition, the use of these simulators requires prior knowledge about the shape and size of the objects, or recurrently testing a set of shapes and sizes [11]. Even with known models, the problem is how to consider each part of the object, including any shadows or overlaps that may be made between them, in order to correctly simulate the measurements.

This paper presents a methodology that seeks to overcome the aforementioned problems, not by reinventing or breaking with established techniques, but by employing a multiple-step approach, chaining together algorithms, in some cases with particular innovations, to synthesize information from diverse sources. The scope of the study directly addresses characterisation in terms of estimating the size and attitude or attitude mode of objects, including determination of rotation if applicable, but the methodology can also be applied to approximate their shape or materials. The main types of data used are light-curves and RCS, except for attitude estimation, for which the consolidated study has been carried out with light-curves only, with the use of RCS currently under development. In the case of attitude estimation, due to the need for a simulator, the development of a light-curve simulator started some time ago at GMV, based on OpenGL, has been concluded and now provides a very valid approximation taking into account the different parts of the object and their interactions. This paper introduces this simulator and explains its main features, although it is not its main goal.

2. METHODOLOGY

This section is dedicated to present the proposed RSO characterization methodology, indicating each step with its logic, justification, the elements to be determined and the data fusion process in each one according to its inputs.
As mentioned above, this is a multi-step methodology to chain several algorithms so that a deeper degree of knowledge of the object is gained after each one and data fusion is exploited to the maximum. The main reason for proceeding in this way is the major difficulty of directly attempting attitude estimation when the level of prior knowledge of the object is very low.
A high-level schematic of the process is shown in Fig. 1, which is explained in the paragraphs below.

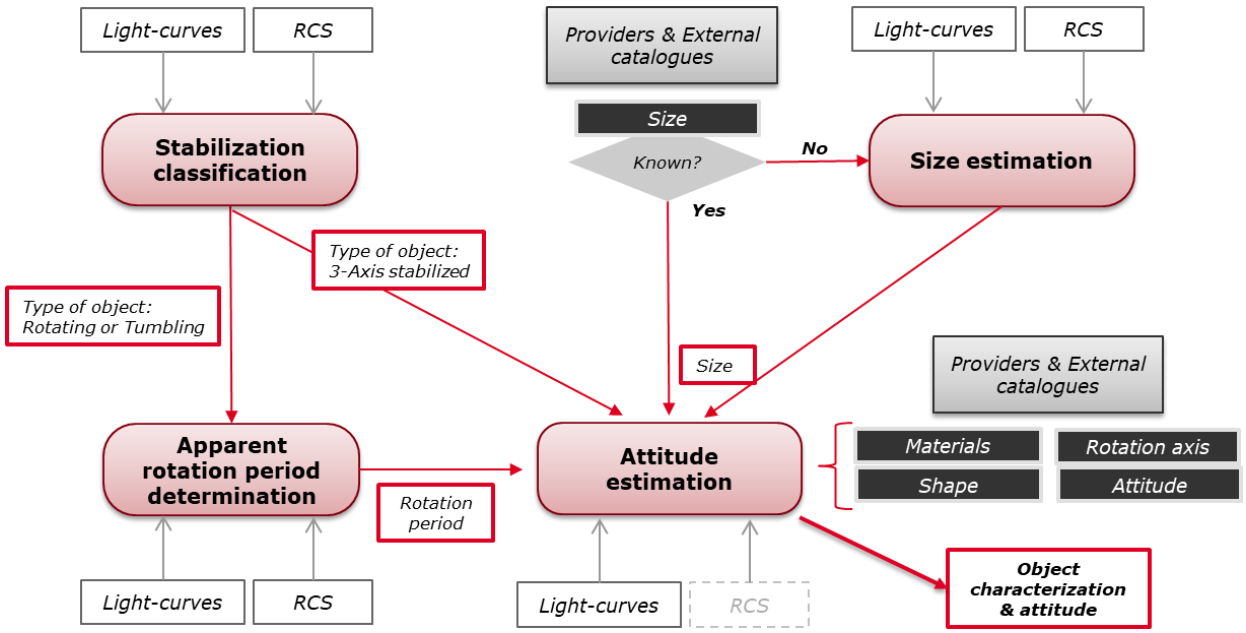

Fig. 1. Schematic representation of the proposed characterization methodology

For the ultimate goal of estimating the attitude mode and the main attitude parameters, they will be very different depending on the type of object stabilization, and therefore a prior classification is essential. The process starts (top left in Fig. 1) with this classification, being necessary to differentiate between 3-axis stabilized objects, thus following an attitude law, and tumbling or rotating objects. As explained in Section 2.2, along with other details, this process is done with Machine Learning (ML) and based on light-curves and/or RCS. The result of this step is just the identification of the object under its corresponding classification.

While 3-axis stabilized objects could pass directly to attitude estimation, tumbling objects are first processed by another step to determine their apparent rotation period (bottom left in Fig. 1). Although the rotation period of an object could be estimated as another parameter in the attitude estimation, it has been observed that the error obtained when there is a lot of uncertainty in its initial value is unacceptable high. However, as it is shown below, the attitude estimation of rotating objects (e.g. the inclination of the axis) is much better when the period is known in advance with a higher accuracy. Moreover, the proposed algorithm, based on Lomb-Scargle periodograms with refinement by Epoch Folding (see Section 2.3), has been shown to give very good results. This algorithm uses light-curves and RCS as input. Although the main output is the apparent rotation period, it is important to note that this step can sometimes serve to cross-check the result of the previous one, as this method can identify objects that have been previously classified as "rotating or tumbling" which are actually not, thus solving possible errors in the ML and providing a more robust solution to the methodology.

Before the attitude estimation, there is one parameter very relevant to know: the size of the object. While a scale factor (see Section 2.5) can be estimated during attitude computation to consider this effect to some extent, it has been found that the results obtained are significantly better when the initial size is known (note that light-curves and RCS are directly impacted by the objects size). Therefore, another step in the chain is to examine whether there are external

sources (catalogues or sensor information) that reliably provide an a-priori value for the size of the object. If not, a new algorithm (top right in Fig. 1) is in charge of estimating it based, again, on light-curves and/or RCS (see Section 2.4).

Once at this point, the next step is the attitude estimation (bottom center in Fig. 1), in this case using only light-curves. The major innovation in this methodology is found using an iterative LSM filter, not commonly used for these purposes. It is used together with a light-curve simulator based on OpenGL, GRIAL, (see Sections 2.5 and 2.5.1), which has been proved to provide good results. It is at this point that the largest representation of the data fusion comes together, as in addition to the light-curves, it also includes the use of various types of a priori knowledge of the object from different sources, namely: size (when available beforehand), shape, materials and initial estimates of rotation or attitude. When this information is known, it becomes part of the 3D modelling required for GRIAL. On the other hand, when this data is not available, a multi-model approach can be used, whereby different typical combinations are tested with standard attitude laws, so that it can be determined which one gives the best results and, therefore, is more likely to represent the real solution. Finally, it should be noted that the GRIAL simulator is currently being improved to simulate RCS measurements, and therefore, it is expected that soon this attitude estimation algorithm will also be applicable to this type of measurement.

Thus, the following subsections present each of the mentioned algorithms with their corresponding details, starting with a brief explanation of the data fusion concept used. Once each step has been understood, the last subsection summarizes the main facts regarding their joint chained use and the cross-checks performed and explains the scope of the methodology as a whole.

2.1 Data sources and data fusion

The previous paragraphs and Fig. 1 have indicated the use of the different data that are relevant in the methodology, without detailing the sources. In the Introduction of this paper, it has been already discussed that telescopes and radars often provide visual magnitude and RCS on their observations. Both types of data are considered in the proposed methodology. However, the use of direct characterization data from *Providers and External catalogues* has also been mentioned and shown in Fig. 1, although its origin has not yet been indicated. On the one hand, optical sensors and radars reappear here, especially the more advanced ones, capable of providing size, attitude, and rotation rate estimates. In addition, some optical sensor technologies can provide accurate information on the shape and materials of objects, although these cases often require manual processing. Using such data has the advantage of increasing autonomy, as there is no need to rely on other sources. The main problem with the use of this information is that, contrarily to light-curves or RCS, much more extended (even considered in the CCSDS Tracking Data Message [5]), there is no standard format for it, which implies making ad-hoc developments or manual configurations for the processing of this information. On the other hand, there are some frequently used public catalogues: European Space Agency (ESA)'s DISCOS Database [6] (with information such as shape and size), Celestrak [7] (with size information) and Space-track.org [8] (with size indications according to categories). In addition, these catalogues also provide information about the type of object, which can be indirectly helpful, for example, by being able to make assumptions such as that "Rocket body" objects are cylindrical in shape even if this information is not explicitly given. Of course, the advantage of being public and free catalogues also implies the disadvantage that they do not contain information about non-public objects.

The combined use of these sources and their data in the proposed methodology exploits the concept of Data fusion at all levels. In their multiple stages, the algorithms are designed to use "the data" (e.g. size, series of RCS data…), independent of the source, and in many cases in first approximation are even independent of the orbit of the object. This allows to refer to data fusion by mixing different types of data, different sources, or simply several data of the same type and source but at different times. As the availability and realisability of data is not always as it would be desired, being able to make the most of the available data is what gives power to the methodology described here.

2.2 Object classification by stabilization type

This first step to provide the binary classification (stable or rotating objects) is carried out using a Machine Learning (ML) method. It extends a previous work developed in [12] for RCS data but considering now the use of light-curve data sequences. Light-curves provide a specific challenge for the ML pipeline, as they are of arbitrary length. Thus, a recurrent-based architecture has been chosen to handle variations in the number of observations in a single track.

2.2.1 Model architecture

Information flows forward through the model as shown in Fig. 2. As observations are fed into the model, they first encounter a Long-Short Term Memory (LSTM) layer [13]. This layer is the recurrent part of the network that allows to process variable sized inputs. An LSTM layer is very useful in this case because it can handle very long sequences of observations without suffering from vanishing or exploding gradients during training. Then, information goes through a two-layer dense neural network which formats and generates a prediction.

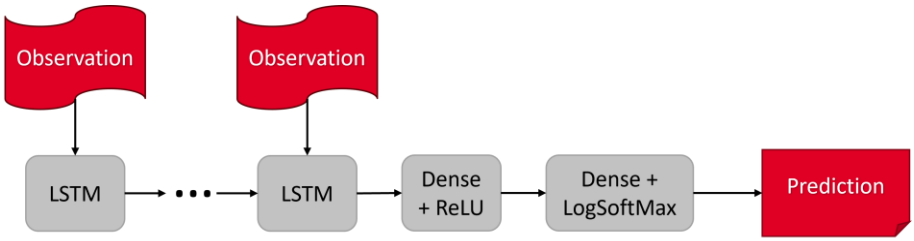

Fig. 2 Machine learning model architecture

Both layers in the feed-forward network have non-linear activation functions. The first one is the usual Rectified Linear Unit (ReLU) ($y = x$ when $x > 0$ and $y = 0$ otherwise) and the output is processed by a *LogSoftMax* function. This function is a logarithm over the output of a *Softmax* function. A *Softmax* transforms a vector of values such that the sum of the elements is equal to 1 using the formula in Eq. 1:

$$Softmax(\vec{x}) = \frac{e^{x_i}}{\sum_{j=1}^{K} e^{x_j}} \tag{1}$$

Each element in the output vector of the formula represents the percent chance of the observation belonging to a particular class. For instance, if the resulting vector is $[0.2, 0.8]$, the sample has a 20% chance of belonging to class 1 and a 80% chance of belonging to class 2.

2.2.2 Training methodology

The training methods used are very important to maximize the performance of the resulting models and data. First, the model architecture shown in Fig. 2 does not show dropout layers that are included in between every layer in the network when training. Dropout layers set to 0 the output of a random subset of outputs of a layer. This results in a more stable training that is less prone to overfitting to the training dataset. The proportion of dropped outputs is a hyper-parameter to be tuned.

Then, the usual ML training process is followed, where the training set is fed through the model in batches, a batch loss is computed, and backpropagation takes place to update the weights in the network. However, a detail is included here to increase the variety seen by the model in training. At every epoch a random set of tracks of the same object are concatenated together before being fed into the model. This brings two main advantages. As it was already said, the apparent entropy of the dataset increases, which leads to a better and more complex model being achievable. Additionally, the model learns to handle several tracks at the same time which is the realistic scenario to be found where more than one track is expected. This has the additional advantage of taking care of the data fusion within the model itself as the model learns to fuse several sources to produce a singular output prediction.

The training is stopped by an early stopping mechanism. The training itself is carried out on the training set but simultaneously performance is measured on a separate validation set. The training is considered concluded when the performance on the validation set no longer improves.

2.2.3 Dataset description

One critical element when it comes to final model performance is the quality and size of the dataset. The light-curve data used to train models is sourced from the MMT9 telescope [14]. It comprises all the tracks available for 214 individual objects chosen arbitrarily. The data is broken down into individual tracks each containing a sequence of magnitude values at corresponding epochs.

To use this dataset for ML it must first be divided into training, validation, and testing subsets, taking great care to ensure datasets are independent and balanced. The target sizes of each dataset for training, validation and testing are 60%, 20% and 20%, respectively, over the total tracks. For them to be independent, if an individual object's tracks are present in one dataset, the other sets must not contain any track of that object. Additionally, the distribution of

objects in the subsets is done such that there are approximately the same number of stabilized and rotating tracks in the validation and test sets. This is done to ensure the accuracy of the metrics used to evaluate the model performance. The labels for whether an object is rotating or not are gathered from the MMT9 dataset directly in most cases (although it has been verified in some cases of which there was prior knowledge). This can cause some issues as the dataset generates these labels directly from its own tracks. However, by using many objects, in the aggregate misclassifications are ignored by the model as in general this classification is accurate.

2.3 Apparent rotation period determination

The problem of finding a period of RSOs using observations has been widely studied, often adopting techniques already used in astronomy (e.g. [15], [16], [17]). Most of the existing algorithms to obtain the apparent rotation period from sensors measurements are based on the Lomb-Scargle periodogram [18][19] analysis (see Section 2.3.1), also known as Spectral least squares analysis [20], which provides the main frequencies of a signal. The Lomb-Scargle is chosen as it is robust against sparse, and unevenly sampled data, which is often the case for the tracks provided by sensors, thus it is often used as a first estimate of object rotation properties.

Another common technique is the so-called Epoch folding [21], which consists of *folding* a track at given periods to find the best fit (see Section 2.3.2). Epoch folding is often used to find the period amongst some previously found candidates, and it can also be used to refine the search once the best period has been found. Additional examples of the usage of epoch folding can be found in [22].

The approach chosen for the proposed methodology employs both the Lomb-Scargle and the Epoch folding techniques to find the apparent rotation period of an object. The Lomb-Scargle periodogram analysis identifies the main frequencies, which allows to have candidate rotation periods. These candidates are filtered and are then used in an Epoch folding process to find the best refined period. One of the strong points is the opportunity to combine data from not only different sensors but also from two different types of sensors, i.e. light-curves and RCS from telescopes and radars respectively. It could also be applied to laser ranging measurements, which is planned to be studied in a future work. Even if there is only data available from one sensor, several tracks can be fed into the algorithm to obtain a more reliable result, making the algorithm more robust against low-quality tracks.

2.3.1 Lomb-Scargle periodogram

As previously stated, to obtain the period candidates, first the Lomb-Scargle periodogram is obtained for each track available. However, for light-curves, the apparent magnitude is transformed to irradiance values (see equation 2), since they provide better results with the Lomb-Scargle algorithm.

$$I_x = I_{x,0} \cdot 10^{-\frac{m_x}{2.5}} \tag{2}$$

where I_x is the irradiance to be found, $I_{x,0}$ is the reference flux (from the Vega star) and m_x is the apparent magnitude.

The RSOs measurements often display a parabola-like trend due to the movement of the object with respect to the sensor. To account for this, the second order polynomial fit has been subtracted from the measurements to flatten the track, thus improving the results of the Lomb-Scargle. Fig. 3 helps to understand the process followed.

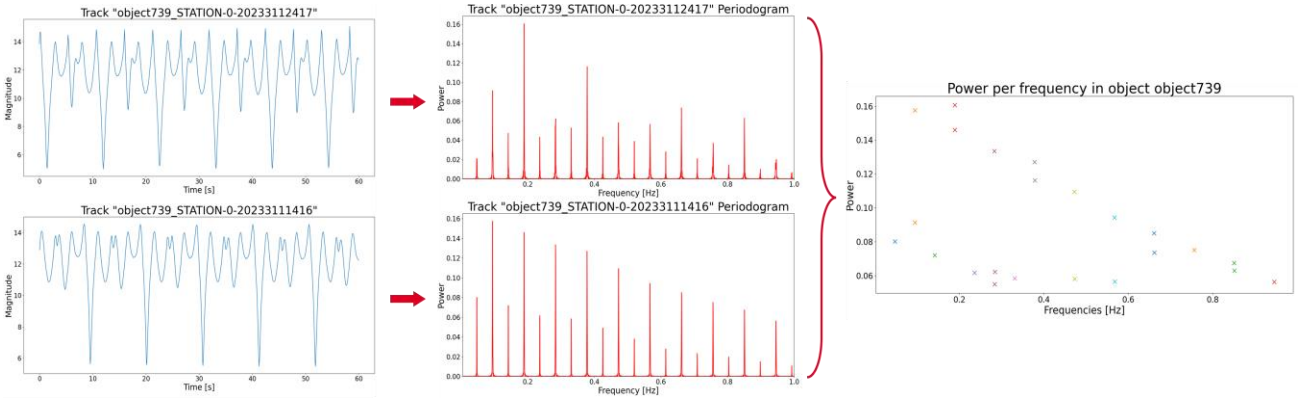

Fig. 3 From left to right: light-curve from two different input tracks; periodograms extracted from the tracks; and frequency peaks found after a first filter (one *color* for each *frequency family*)

For each periodogram obtained, the resulting frequency peaks are stored. Then, peaks from all tracks are grouped by *families* according to their frequency value, i.e. frequency peaks with similar frequency values are grouped within the same *family*. Once these families are found, they undergo two filters to eliminate spurious candidates. The first filter only keeps the frequencies that are found a minimum number of times per track, improving the robustness of the algorithm when using several tracks from the same object by eliminating false frequencies. The second filter removes the frequencies that do not have a minimum of accumulated power, i.e., the sum of powers within a frequency must reach a minimum threshold. These filters are introduced as part of the proposed methodology, and they have proven to improve both the success rate and the computational performance of the algorithm.
Most of the target objects have geometric shapes with symmetries such as cylinders and parallelepipeds, which imply that the rotation period to be found must be within a sequence of multiples within the frequency candidates that passed the two previous filters. Hence, the frequency candidates that are within a harmonic sequence move onto the next Epoch folding phase.

2.3.2 Epoch Folding

The epoch folding technique is used here not only to find the rotation period amongst the remaining candidates, but also to refine it. First, it is remarked that only light-curves measurements have been used for this phase, given that no noticeable improvement has been seen when using RCS real data. For the candidates, each light-curve track available is folded, with the folding done normalizing the magnitude from 0 to 1 in each period. Fig. 4 shows the same track folded at two candidate periods that passed all the filters, however it can be clearly observed than one is a much better fit (right) than the other (left).

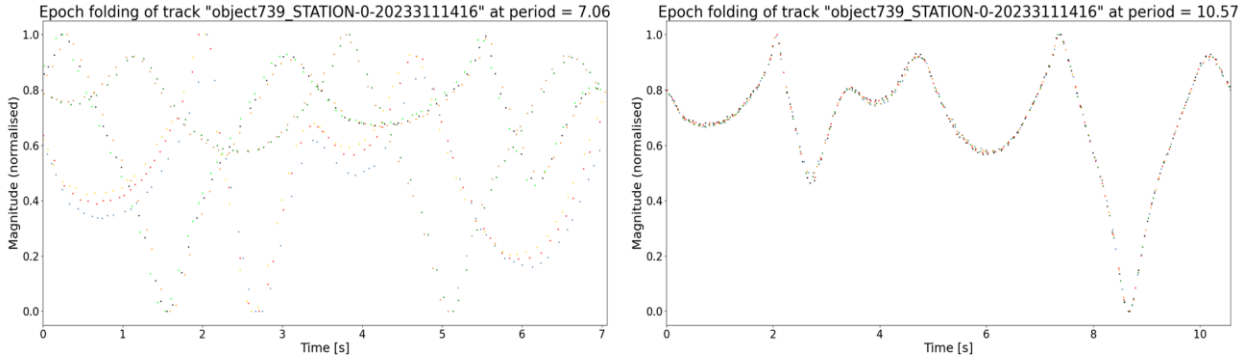

Fig. 4. Epoch folding of the same track at two different periods, 7.06 (left) and 10.57 s (right). Different colors denote the different track sections folded

A method has been developed to numerically evaluate the good match of a folding, so that the best period can be automatically selected. Once the track is folded, it is divided into segments, and for each segment, a polynomial fitting of second order is found. From these parabolas, several points are stored: the middle point of each one, the middle-points between them, and the first and last point from the first and last parabolas, respectively. These points are then used in a spline cubic interpolation to create a continuous line, with that line as reference, the Standard Deviation (STD) with respect to that line is computed. The better fit the folding performs, the lower the STD will be, since the actual measurements will be closer to the spline interpolation. Fig. 5 illustrates this process.

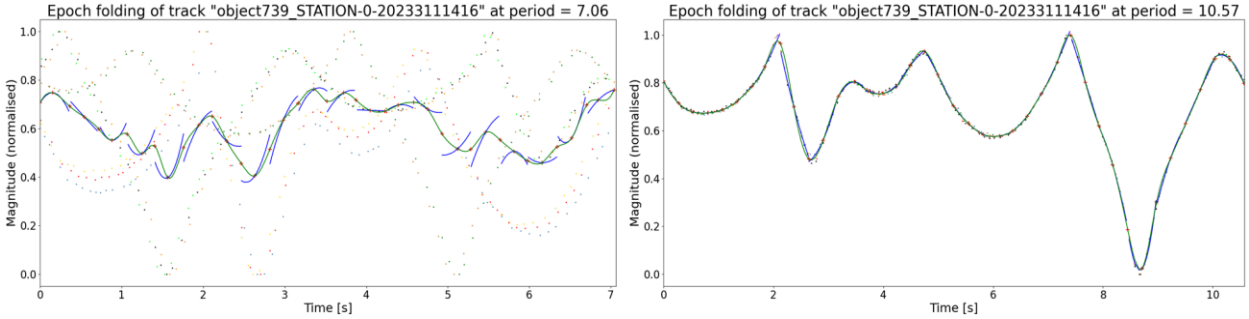

Fig. 5. Same epoch folding as in Fig. 4 with the second order polynomial fitting per segment (blue parabolas), the points extracted from these parabolas (red crosses) and the spline interpolation (green line).

Only one track can be used per Epoch folding, so to benefit from the availability of several tracks (as in Section 2.3.1) for a given period, the STDs of each track are added and then divided by the number of tracks. This resulting number is used to evaluate how good the folding is, i.e. how accurate the tested period really is. For example, in Fig. 5 right, the mean of the STDs is 0.013, and for the one on the left is 0.210.
Now that a candidate period can be numerically evaluated, every candidate period found from the process described in Section 2.3.1 is refined and compared with the others. The refinement process is done with the Nelder-Mead method [23], which allows to find the minimum of the function that computes the mean of the STDs at a given period. Hence, the periods compared have already been refined.

2.4 Size estimation

The size of an object is difficult to estimate with sensor measurements. Nonetheless, estimation models have been studied for both light-curves [24] and RCS measurements [25], which have been implemented here. Note that for both cases the model assumes a spherical shape for the object.

2.4.1 Size estimation with light-curves

The spherical assumption implies that there is no oscillation of photometric measurements, as a sphere presents the same return in all its rotational orientations, and its apparent magnitude (M_v) is given by equations 3, 4 and 5 [24] below:

$$M_v = M_{Sun} - 2.5\log(A\rho[\beta F_1(\phi) + (1-\beta)F_2(\phi)]) + 5\log R \tag{3}$$

$$F_1(\phi) = \frac{2}{3\pi^2}[(\pi-\phi)\cos\phi + \sin\phi] \tag{4}$$

$$F_2(\phi) = \frac{1}{4\pi} \tag{5}$$

where M_{Sun} is the Sun's apparent magnitude (-26.74), A is the cross-section area of the observed object, ρ is the albedo of the observed object, β is the diffuseness of the object (1 being purely diffuse and 0 purely specular), ϕ is the phase angle (angle between the object-sensor vector and the object-sun vector) and R is the distance between the sensor and the object.
If all parameters are known or estimated, one can compute the size A of the observed object. The albedo and the diffuseness are parameters that depend on the materials and the surface of the object, and both can go from 0 to 1. Typical albedo values are found in [26].
Each track contains different measurements, and for each measurement, a size is calculated. Then the median is used to account for outliers that may occur due to brightness peaks at specular reflections or when the sensor is receiving light from darker areas of the object. If more than one track is available, the mean size with all tracks is calculated.

2.4.2 Size estimation with RCS

The size estimation with RCS has been implemented with a Size Estimation Model (SEM) developed by NASA [27], which is based on static laboratory RCS measurements for a set of thirty-nine representative debris objects. The SEM (see the formulation in Section 3 within [25]) has been implemented to find the estimated size for each RCS measurement.
As when using light-curves (see 2.4.1 above), for each track the median of the resulting size values is taken, and, if several tracks are available, the mean of all tracks is then computed.

2.5 Attitude estimation

The final stage of the whole process defined in Fig. 1 is the attitude determination. This process obtains the complete attitude ephemeris of the observed object in the periods where magnitude observations are available. As this is the last stage, it gives the most detailed information on the attitude of the RSO, however, it also requires more prior knowledge, which limits its application.
The underlying method used is a Least Squares Method (LSM) filter that relies on a light-curve simulator to generate estimated observations based on a set of determined parameters that define the attitude of the object. In this section, each of the elements that make up this process are explained except for the light-curve simulator, which is discussed in Section 2.5.1.

LSM filters [30][31] are algorithms that allow to the minimize the squared error on a signal of another estimated signal by finding a set of optimum parameters. The filter is iterative and explores the non-linear function that relates the estimated signal with its parameters by performing a gradient descent on the squared error. This process is very flexible and allows the estimation of individual parameters and locks others depending on prior knowledge of the problem. Indeed, this is the basis for many of the most widely used orbit determination algorithms.
However, generating a complete attitude ephemeris is problematic without providing some additional constraints on the system. Therefore, a set of assumptions are considered. First of all, certain prior knowledge of the object is assumed as known. That is, its shape, size, and materials (for their light reflecting properties) must be known to a certain degree. Additionally, the object is presumed to follow a constant attitude mode, as is most often the case for RSOs. Several possible attitude modes are considered:

- Inertial pointing: The object keeps a constant attitude with respect to an inertial reference frame. This attitude mode is common for space science missions, like telescopes. The parameters required to be estimated are the elements of a quaternion that defines the rotation with respect to an inertial frame.
- Earth pointing: The object keeps one axis (+Z by convention) pointed to ground and the axis perpendicular to the solar panels (+X by convention) pointed to the sun as much as possible. This attitude mode is typical of Earth observation missions. This attitude is defined as a bias rotation over the nominal attitude to account for possible misalignments between the assumed pointed axes and the actual attitude.
- Earth pointing with yaw compensation: Similar to straight Earth pointing but with an additional yaw rotation (around Z) to reduce distortion when imaging the ground. Typical of larger Earth observation missions.
- Safe pointing: The spacecraft points its panels to the sun directly (+X axis). This can be considered an inertial pointing in the time frames considered for this application. It needs another parameter to define the rotation around the sun pointing axis which is assumed fixed during a single track.
- Spin stabilized: The spacecraft maintains a constant spin around an inertial axis. This attitude mode can be defined with an orientation at an epoch plus a spin axis and spin rate, that is 7 total determination parameters. However, an estimate of the spin rate has been estimated before with the methods in Section 2.3 (as the approximation with the apparent rotation period is considered to be enough in most cases). And the spin axis is typically known (from the shape of the object) for stabilized objects and simple shapes like rocket bodies. For tumbling bodies, a first estimation of this axis can be found with for LEO objects with Inverse Synthetic Aperture Radar (ISAR) or adaptive optics sensors.

With these attitude modes, most active satellites can be covered. However, tumbling debris still pose a challenge as they typically do not follow what has been defined as a spin stabilized attitude. Instead, their attitude is much more complex with precession and nutation that is affected by external torques. Nevertheless, for short time frames, the spin stabilized attitude can approximate this attitude, if not completely, enough to give an idea of the attitude of the object.

Simulating the observed magnitude poses certain difficulties in that there is strong coupling between the optical properties and the size of the RSO. For instance, a very reflective small object may have a similar luminosity to a large but dull object. Additionally, optical parameters, not only are typically not well known, but they may change over the life of the satellite. Because of this coupling, when uncertainty is large in this regard, a scale factor parameter can be estimated. This scale factor effectively changes the size of the spacecraft. Estimating this parameter must be done carefully as it can interact with the attitude such that the LSM filter finds a spurious local minimum. Additionally, atmospheric parameters at the time of the observation are required. However, these are rarely measured and so typical values that have been shown to work well in practice are used. Of particular importance to this methodology is the Aerosol Optical Depth (AOD) which is fixed at 0.15.
Having described all the parts that make up the system, a typical execution goes as follows. A simplified 3D model of the object is created. When available, it comes from images or detailed plans and based on the material of its external surfaces, specular and diffuse coefficients are assigned to each part. Then, a LSM determination process is carried out to try to fit each of the attitude modes sequentially to the observations analyzed. This will yield an attitude ephemeris of a best fit for each attitude mode and an associated Root Mean Squared (RMS) error on the true measurements from the measurements simulated with that attitude mode. Finally, based on this RMS, a most likely attitude is selected as the actual attitude being observed.

Additionally, it is worth noting that the flexibility of this methodology allows for other uses that may also be interesting. As it was mentioned earlier, optical properties of the external surfaces of the spacecraft change over time. It is therefore useful for satellite operators to keep this information up to date in case a loss of control of the satellite. This can be achieved by providing a detailed satellite model and accurate attitude ephemeris. In this way, the optical

properties can be released as parameters safely of couplings to be determined. Then, in the event of a communication loss with the object, an attitude determination can be performed that would be that much more accurate.

A final important remark is that, as the LSM filter is iterative and relies on derivatives on each of the parameters at each iteration and observation, this poses strong requirements on the light-curve simulator to be very computationally efficient. The derivatives with respect to the attitude related parameters cannot be obtained directly and must be obtained by numerical differentiation. This means that for every iteration and observation, for N released parameters, the simulator must perform N+1 evaluations. Moreover, it must be accurate to properly capture common optical effects such as flares and self-shadowing, where a piece of the spacecraft covers others. The approach for the simulator is discussed in the following section.

2.5.1 GRIAL: Light-curves simulator

GMV's GRIAL (Generator of RealistIc shApe/Attitude reLated data) is a high-accuracy, high-performance light-curve simulation tool developed in house to tackle among other challenges, the attitude determination process that is discussed in this paper.

It considers a Bidirectional Reflectance Distribution Function (BRDF) over a 3D shape, with a given attitude and materials to compute an apparent magnitude at an observer. This BRDF provides the ratio of reflected light in a given direction, considering a specular and a diffuse term, discarding the usually present ambient term which is not appropriate in space. The specular term is the mirror-like reflection off a surface, characteristic of smooth surfaces. The selected BRDF function for this term is the one proposed by Cook and Torrance in [32], but with the Beckmann's distribution proposed there replaced by the one proposed by Walter in the GGX model [33]. On the other hand, diffuse reflections are those that do not reflect in a single direction, instead they scatter in many directions off the surface. It is characteristic of rough surfaces. The BRDF function chosen for the diffuse term is the one from Lommel-Seeliger [34], however, it has been modified to better respect the conservation of energy and consider a changing diffuse reflection coefficient with reflection geometry.

In these terms, the effects of self-shadowing must be also considered. This occurs when an illuminated part cannot be seen by the sensor because other part of the object covers it or when a part visible by the sensor is not illuminated because another one covers its view of the sun. These effects are typically computationally expensive, however, GRIAL is implemented using OpenGL [35], a cross-platform programming interface for 2D and 3D graphics. Its use allows for these calculations to be carried out in the Graphics Processing Unit (GPU) of a computer at much faster speeds than would be possible in the Central Processing Unit (CPU).

All of this allows to compute the reflected radiance of the object in a direction, however some effects must be considered to compute the magnitude at the station. The incident irradiance is the radiant power from the Sun that is received by the object. It decreases by an inverse square law of the distance between the Sun and the object. Additionally, partial, or total Sun eclipses of the Earth are considered to compute the total incident irradiance received. Similarly, after light is reflected off the object, the observed irradiance at the station decreases with the square of the inverse of the distance from the object to the station. Besides attenuation due to the distance, the atmosphere also plays a role, using the attenuation model defined in [36].

Care must also be taken when considering the geometry of the light path. The reflection date must be computed considering the speed of light as well as the delay caused by the troposphere. This last one is modelled using a Mendes-Pavis model [37] which is well adapted to optical measurements. Finally, with a computed observed irradiance at the station, one can obtain the apparent magnitude as follows:

$$m = -2.5\log\left(\frac{I_o}{I_{ref}}\right) \tag{6}$$

where I_{ref} is the reference irradiance corresponding to the zero point. As the incident irradiance of Sun has been computed from the luminosity of Sun in the whole electromagnetic spectrum, the reference irradiance here is the one corresponding to the zero point of the bolometric apparent magnitude: $I_{ref} = 2.518 \cdot 10^{-8} \frac{\mathrm{W}}{\mathrm{m}^2}$.

This methodology allows accurate and high-speed magnitude simulations that make iterative attitude determination possible in reasonable time frames. Performance and accuracy are discussed in Section 3.4.

3. RESULTS

The following subsections contain a summary and a discussion of the main results obtained with each of the algorithms presented in the previous section, according to a variety of real and simulated cases, which help to better understand the scope and use of the proposed methodology.

3.1 Object classification by stabilization type

First, the results obtained in attitude stability classification by the ML models trained (see Section 2.2) are covered. To do this, some useful metrics must be defined. Since the classification that is performed in this case is simply binary (an object can either be rotating or not) the results are given in terms of a 2x2 confusion matrix. It allows us to interpret the accuracy, precision and recall of a model on a dataset at a glance, by providing the fraction of correctly and incorrectly classified positive and negative samples (Fig. 6).

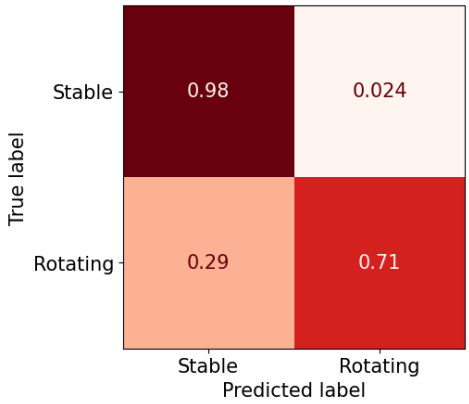

Fig. 6 Confusion matrix using 5 tracks as input.

The confusion matrix shown in Fig. 6 has been computed on the test dataset defined in Section 2.2.3, using 5 tracks from the MMT9 telescope for each object. Each object is considered a sample in this case. The number in each box is the proportion of the true labels in each category, therefore rows add up to 1. It is clear from the results that the model can accurately classify objects into stable or rotating, in fact, the overall accuracy in this case was 83%. However, the model is biased toward stable object as 30% of actually rotating objects are classified as stable. This bias is probably due to imbalances in the dataset and can be mitigated by introducing a bias in the opposite direction in post-processing. The plot showed the results when 5 tracks are available for each object, however, it is interesting to explore how performance changes as a function of the number of tracks used.

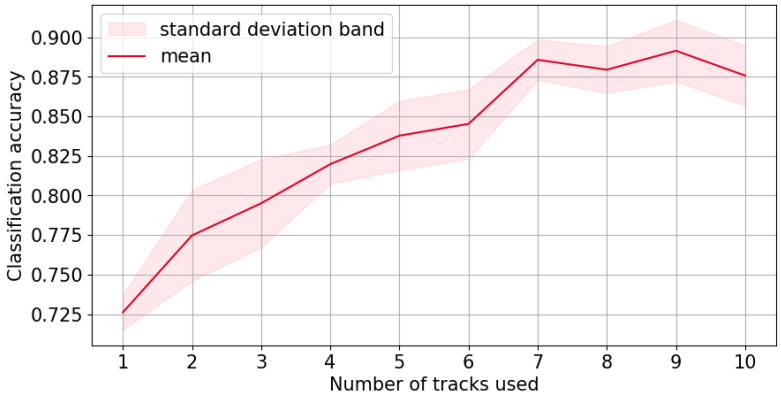

Fig. 7 Model accuracy as a function of number of tracks per object

As expected, the model yields better results when more tracks are available. From Fig. 7, it is clear that the performance of the model increases steadily from ~75% with only one track up to ~90% using ten tracks. To generate this plot, for each object, a random subset of its tracks available in the test set was drawn ten times so that the variability of the

performance can also be presented in the form of a standard deviation around a mean accuracy. To achieve this result, the model is performing data fusion on the separate tracks in order to gain more in-depth knowledge of the observation than would be possible using single tracks.
This result is consistent with what was observed using RCS from radar tracks in [12], with those results achieving a max accuracy of ~95%. These small differences may be attributed to differences in the available dataset and the breakdown of training and testing sets.
Finally, it is important to note some limitations of this algorithm. A very poor performance was observed when using real data from other telescopes. This is expected because the model was trained solely with data from MMT9. In the future, if this algorithm is to be deployed to analyze tracks from a sensor network, the training dataset must consider tracks from a variety of telescopes with different characteristics such that the model becomes general over the individual sensor used.

3.2 Apparent rotation period determination

Then the validation of the algorithm described in Section 2.3 can be discussed according to the results obtained.
The algorithm has been firstly tested with a small set of simulated tracks from GRIAL, containing a total of 75 different rotating objects and several tracks per object. Since the exact rotation period is known, the algorithm can be evaluated. The algorithm finds the real period of 49 objects (65.33%), which is less than one would expect when considering that the tracks are noiseless and *ideal*, however, in another 22 objects the period found is a multiple/divisor of the real period. This is caused by the symmetries found in most of the objects in one or more axes (as mentioned in Section 2.3.1) even more pronounced in these simulated cases, due to the use of simplified models. Hence, out of 75 objects, in a total of 71 (94.67%) either the real period has been found or a multiple or divisor. In the cases where the real period has been found, it has been done with a mean error of 0.012%, which shows how accurate it can be under ideal conditions.

To make a more realistic validation, for the case where simulated data may not be fully representative, a larger set of real data has been used. The data used for the validation is from MMT9 [14]. This database contains a large number of tracks and also a rotation period calculated for each track, which allows to perform a direct comparison of the period found. The real data set used contains tracks of a total of 900 objects (with one or more tracks per object).
Out of the 900 objects tested, the algorithm has successfully found the period in 477 objects (53%), and a multiple/divisor of the period in 294 objects (32.67%). This means that only in 14.33% of the cases the algorithm did not find the same rotation period nor a multiple/divisor that MMT9 has calculated. Moreover, within the 477 objects where the period has been successfully found, the mean error is of 0.85%.
Fig. 8 shows several cases with real data, using the same type of plot shown in Section 2.3 to explain the algorithm.

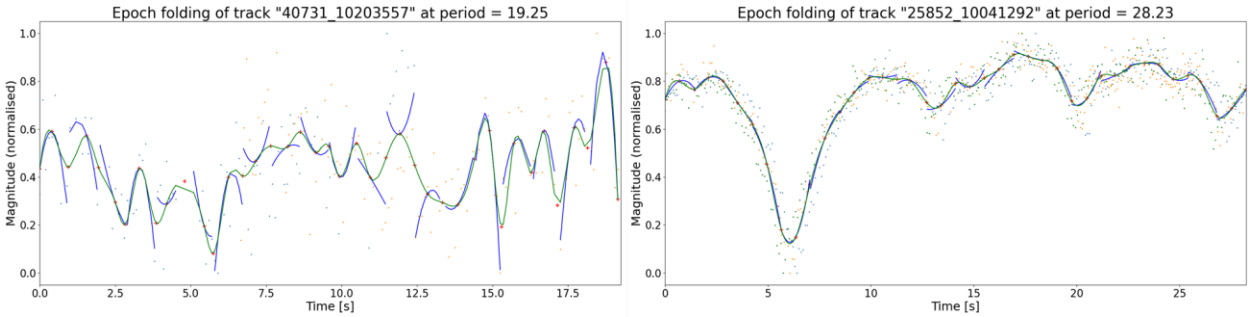

Fig. 8. On the left, an example of epoch folding in a case in which the algorithm failed, on the right, an example of epoch folding in a case in which the algorithm did succeed.

In Fig. 8 left, it is clearly seen that the algorithm did not manage to find the period, mainly because the measurements in this track are scarcer, which creates difficulties to the process (especially in shorter periods). On the right, there is a good example of the algorithm finding the period, with a fitting much better, and all segments of the track following the same trend denoted by the spline line (green).

As a partial conclusion, the proposed methodology has been proven to be reliable in most of the cases in which there is enough quality data. In addition, its accuracy is very high, yielding overall errors lower than 1%. However, more work can be done especially in cases where the tracks have few measurements per unit of time. Also, it is very difficult to distinguish between harmonics, although this can be solved with some manual post-analyses.

3.3 Size estimation

For the size estimation, the validation is not straightforward with real data, since the models used make assumptions that may not represent the reality in many cases.

3.3.1 Size estimation using light-curves

The chosen objects have been spherical ones, given that it is the shape that the model considers. These objects are two geodetic satellites of which real tracks were available from MMT9: STELLA and LAGEOS 1 (see Fig. 9), having a diameter of 0.24m and 0.6m respectively [29].

Using light-curves, the obtained results assuming a default albedo of 0.175 for space debris [26] and a beta value of 1 (i.e. completely diffuse objects) are not correct for any of the two cases. For STELLA, the resulting diameter is 0.48m, double the actual value. For LAGEOS 1, the resulting diameter is 4.95 m, which is even further from the real value. However, for the STELLA satellite, if the albedo value is adjusted (0.7), the correct real diameter is obtained. The satellite surface is made up of polished aluminum and retroreflectors [29], and the albedo of polished aluminum is closer to 1 than to 0.175, hence the obtained results can be considered adequate. On the other hand, even adjusting to a higher albedo value for the LAGEOS 1 satellite, a diameter of around 2m is obtained, still far from the real value. This is most likely because of the light reflectance behavior caused by the retroreflectors, not contemplated in the model.

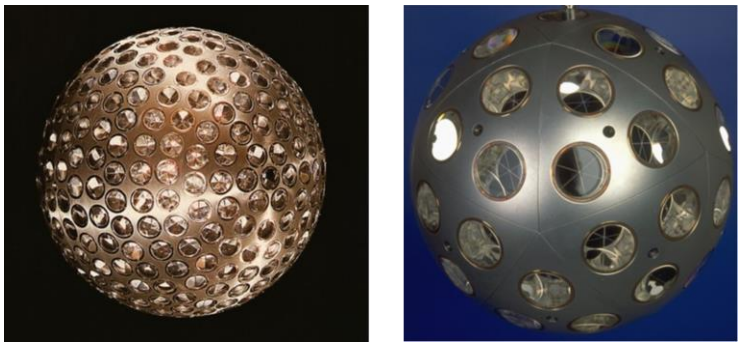

Fig. 9. LAGEOS 1 (left [29]) and STELLA (right [28]) satellites.

The resulting size estimation depends on the parameters of the observed object which are sometimes difficult to know beforehand, thus posing a limitation on the size estimation model. Despite that, getting an estimate of the order of magnitude can still be useful for the attitude estimation, thanks to the consideration of the scale factor (see Section 2.5), especially if no other sources are available.

3.3.2 Size estimation using RCS

As for the previous case, the object chosen to validate the size estimation with RCS data has been a spherical one: COSMOS 660 (diameter of 2m [6]). Using a set of over 100 real radar tracks, the results are closer to those expected, with a mean error of 13% in diameter. These errors could be caused by the model used but also by object observed, which even though it is a sphere, it is not a smooth, perfect sphere.

Even though the result still has an error larger than 10%, it is considered good as initial estimate when no other data is available, or even to cross-check already available data. Moreover, the use of several tracks greatly increases the accuracy and the trustworthiness of the result.

3.4 Light-curves simulation

Before discussing the results of the attitude estimation, the performance of the light-curve simulator is shown and compared against real measurements.

To do this, the satellite SMOS [38] is used. It is chosen out of convenience as attitude ephemeris data is available at the same time as measurements from MMT9. SMOS is an Earth observation satellite that monitors surface soil moisture and is part of ESA's Living Planet Programme. It follows a tilted Earth pointing attitude with its MIRAS instrument pointed ~32º off nadir [38]. To simulate magnitude observations, a simplified model was created. It is shown compared with a more realistic rendition in Fig. 10.

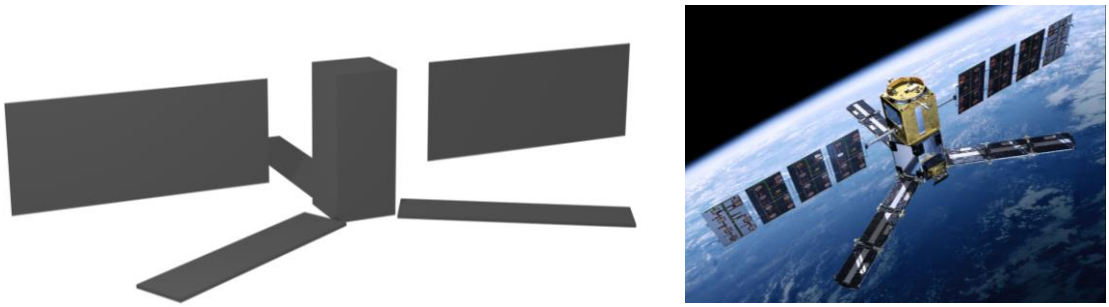

Fig. 10 Simplified (left) and realistic (right [38]) renders of SMOS satellite.

The model is much simpler than the real satellite, however, major reflecting surfaces are included and the proportional importance of each part of the spacecraft is preserved. In addition, assigning the optical properties of each part of the satellite is of importance. In this case, two materials are used: a not very reflective material for the solar panels with specular reflection coefficient of 0.1, diffuse coefficient of 0 and surface roughness of 0.1; and the rest is modelled as a very specular material with specular coefficient of 0.75, the other parameters are identical.

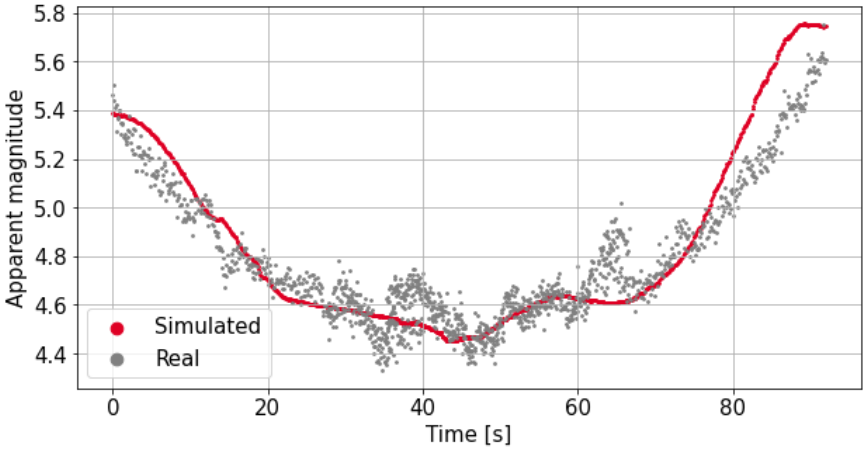

Fig. 11 Real and simulated SMOS magnitude observations.

As is clear from Fig. 11, real and simulated observations agree with each other significantly. The mean error is 0.01 with a standard deviation of 0.12, both in terms of magnitude. Beyond the noise that is not present in the simulated values, it is clear that there are differences in the granularity of the magnitude, as the real light-curve is more detailed. This is a consequence of the lack of detail in the model. Nevertheless, it is useful to see that the main characteristics of the plot are captured with such a simple model.

In terms of performance, GRIAL can simulate observations using an integrated intel UHD Graphics 620 GPU at a rate of ~250 observations per second. This performance allows for the use of the tool within iterative filters which require many evaluations.

3.5 Attitude estimation

Finally, results obtained with the attitude estimation method are discussed. While light-curve data is available from MMT9, real attitude ephemeris data is not. Therefore, to show the accuracy of the algorithm this section is divided in two subsections. In the first one, the algorithm is applied to simulated data with added noise. Results on simulated data are not as valuable. However, it has the advantage that the true attitude of the spacecraft is fully known, and a determined attitude can be compared against it. In the second part, real data is used, and the analysis is focused on the attitude mode determined and magnitude residuals.

In both cases, the satellite used is Sentinel-3A [39]. It is an Earth observation satellite from ESA. It is chosen because it has a known attitude mode and magnitude observations are available. As was done in Section 3.4, a simplified model of the spacecraft must be created, in this case shown in Fig. 12.

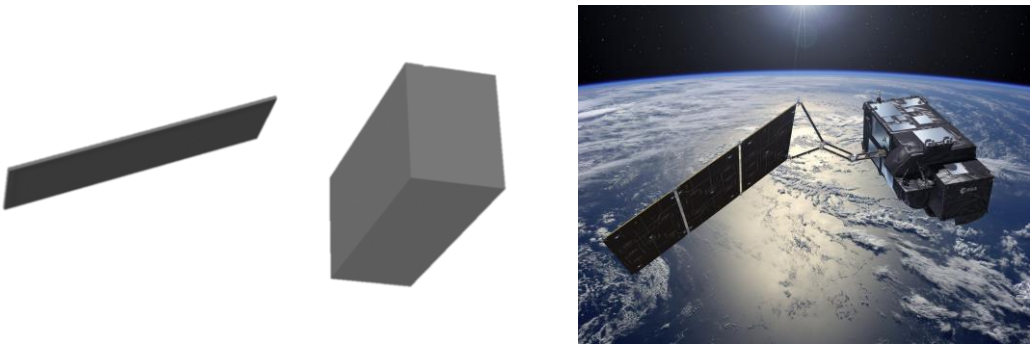

Fig. 12 Simplified (left) and realistic (right [39]) renders of Sentinel-3A satellite.

3.5.1 Results on simulated data

First, the accuracy of the model in finding the true attitude is discussed, together with the sensibility to flawed prior information on the object. This is achieved by simulating two tracks: one with the true size and a second one with a 50% larger model. This simulates a case where the object size is known and another where it is underestimated.

Tracks are simulated with 100 measurements over 5 minutes. To generate the track, the simulated satellite follows a tilted yaw compensation attitude (see section 2.5 for a description of the attitude modes). The raw signal is generated with gaussian noise with mean 0 and standard deviation 0.5. It is assumed that earlier algorithms have already determined that the RSO is not spinning, which discards spin stabilized attitude mode. Additionally, for simplicity, the safe attitude mode is not attempted as it is effectively a particularization of an inertial attitude mode in the time frame of the track used.

Results after running the algorithm with the three remaining attitude modes in the case where the size of the object is initially correctly inserted are shown in Fig. 13.

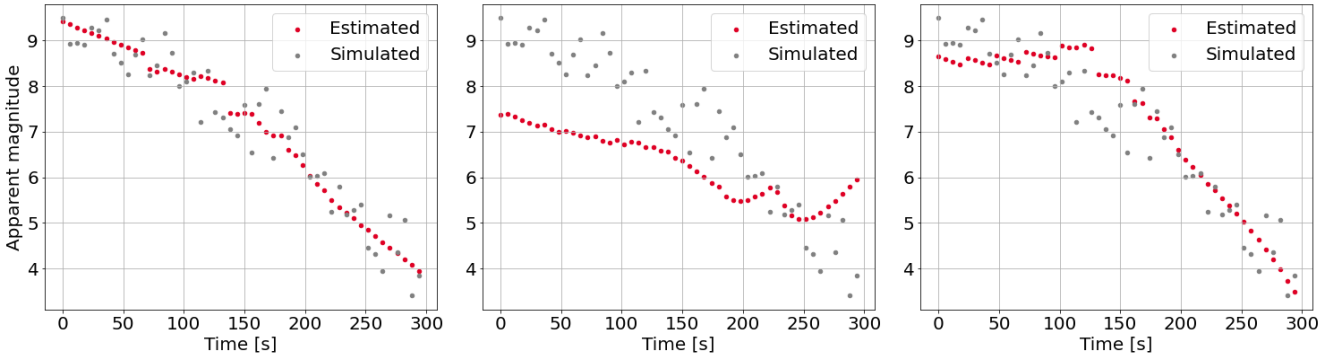

Fig. 13 Simulated and estimated light-curves assuming yaw compensation (left), Earth pointing (center) and inertial pointing (right) with correct size estimation.

From the plots, it is clear that the correct attitude is yaw compensation (left). Indeed, when comparing the weighted RMS error (weighted with the signal noise) of all cases, the results are: 0.87, 2.65 and 1.27 for yaw compensation, Earth pointing and inertial, respectively. In this case, the algorithm automatically selects yaw compensation as the correct attitude mode. Since it is a simulated track, the average angular error of the estimated attitude can be computed, which is in the same order as the plots: 3.3º, 66º and 164.1º. It can be asserted then that the yaw compensation mode determination is very accurate and has completely restituted the attitude of the RSO.

However, it is interesting to notice how incorrect prior information affects the determination process. These results can be seen in Fig. 14.

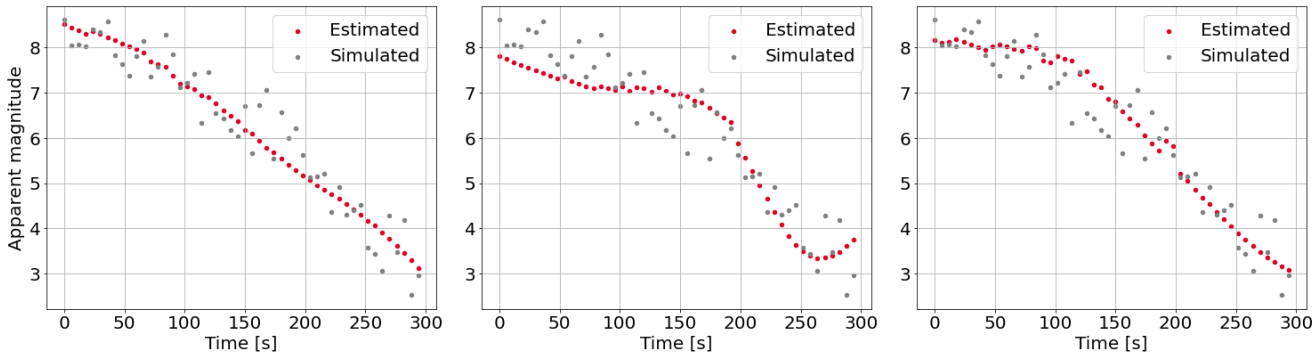

Fig. 14 Simulated and estimated light-curves assuming yaw compensation (left), Earth pointing (center) and inertial pointing (right) with incorrect size estimation.

Again, the resulting RMS error and visual appearance suggest that the object is following a yaw compensation attitude. The RMS in this case is: 0.92, 1.18 and 0.99 in the order of the plots. It is difficult to see meaningful differences in the simulated data looking just at the RMS error. However, the attitude errors in this case are: 49.48º, 40.35º and 155.77º respectively.

Indeed, this issue is in principle intrinsic to the attitude determination process. When an object appears brighter or dimmer, the LSM filter is able to obtain an attitude that reflects more or less light to compensate for the size error on the RSO. This same effect can be observed with the optical properties of the satellite materials and the atmospheric properties around the telescope at the time of the observation.

3.5.2 Results on real data

In this section, some results of the algorithm are tested on real tracks from MMT9. As in Section 3.5.1, results of only three attitude modes is shown as it is assumed that spin stabilized has already been discarded and safe mode is not interesting for the discussion here. The real vs simulated plots are shown in Fig. 15.

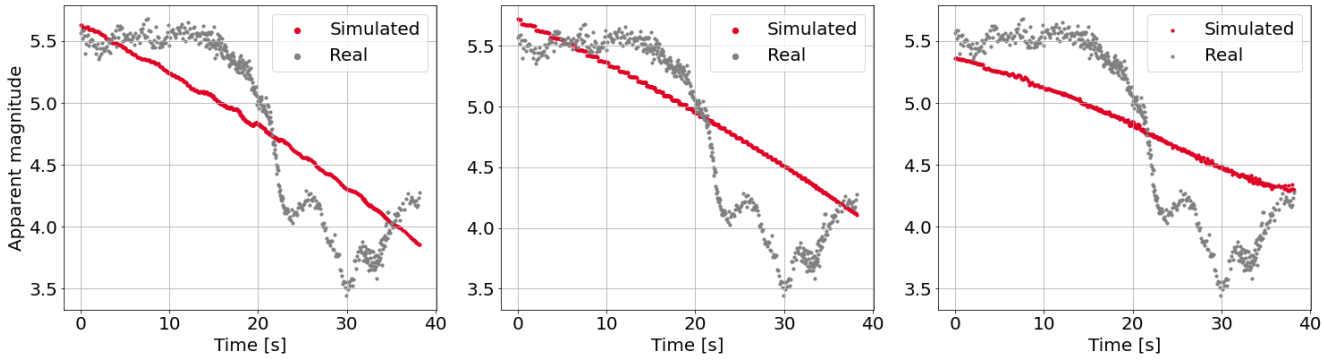

Fig. 15 Simulated and real light-curves assuming yaw compensation (left), Earth pointing (center) and inertial pointing (right) for a real Sentinel-3A track.

The first noticeable thing in these plots is that the simulated light-curve does not show as much variability as the real measurements. As it was already explained in Section 3.4, this is a consequence of the simplicity of the model and, for the purposes of the determination algorithm, it is not critical. Looking at the resulting RMS (in this case unweighted for simplicity), for each case, it is: 0.34, 0.39 and 0.43. The results are very close together even if the process correctly identifies the real attitude mode as yaw compensated nadir pointing. This result is a consequence of the length of the track. For very short time frames, the two nadir pointing modes can be approximated as inertial, since, for an orbital period of ~100 minutes, the deviation in 35 seconds from inertial is ~1.2º. This issue is resolved for longer tracks or multiple observations in short periods of time where data fusion can be applied.

As an additional remark, it is worth noting that GMV participated with this algorithm in the Sprint Advanced Concept Training (SACT) [40] exercises in July 2023, with very promising results when a different source of real optical data was used, being able to completely resolve the attitude of Sentinel-3A as verified by the attitude ephemeris available.

To conclude the explanation of the results with real data, a very different case can be analyzed: an observed object that is not controlled, considering in this case a Falcon 9 upper stage rocket body. MMT9 measurements have been used again, which have previously been run through the apparent rotation period determination algorithm, resulting in a correct identification of the apparent rotation period of 87.5 seconds. Since it is a rocket body, it is modelled as a simple cylinder with the correct proportions as given in DISCOS [6]. Its rotation axis may be initially assumed as pointing perpendicular to the cylinder axis although this may not be the case in general. If we estimate only for the initial orientation of the object, the following result is obtained (Fig. 16).

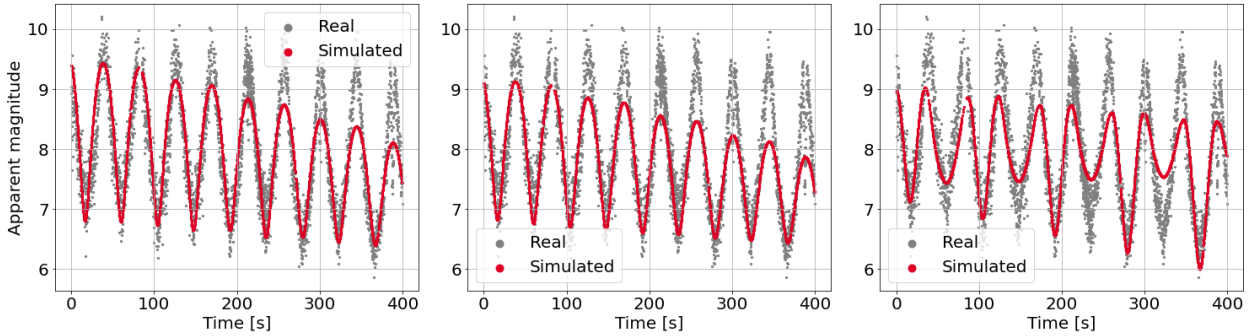

Fig. 16 Spin stabilized determination on Falcon 9 R/B, from left to right: fixing correct rotation axis and rate; estimating rotation rate; estimating rotation axis.

Again, figures match significantly. However, this result can be replicated if the rotation rate or the rotation axis is released as parameter. In that case, the rotation rate is determined with essentially zero error (<4e-4 relative error) and the axis direction has an error of around 11º with respect to the plane perpendicular to the axis of the cylinder. This is consistent with the initial assumption used in the first two determination runs. In this case, the downside is the expense of computational time and a less accurate global estimation. As in previous cases, it can be noticed that a better final result is obtained with a better a priori knowledge of the object.

4. CONCLUSIONS AND FUTURE WORK

Finally, the conclusions drawn from this work are summarised. The future work that can be done is also explained, taking as a starting point the proposed methodology and the results obtained, to overcome current limitations and extend it with new and useful functionality.

In view of the results obtained, it has been shown how the stabilisation of objects can be determined with sufficient accuracy to be able to proceed with more detailed characterization of the RSOs. In addition, the good results in the determination of the apparent rotation period allow it to be used as a preliminary step in the attitude estimation, and it can be used as an a-priori estimate of the rotation rate in most cases. The size estimation, although it provides results that are often not so close to the real ones, is nevertheless enough to have an initial value when there is no other alternative source. Finally, in addition to being a relatively innovative method, attitude estimation has performed well and provided very promising results in a variety of cases, both simulated and real.

Of course, there is still much room for improvement and work ahead. Machine learning methodologies need to be improved with much more extensive training with real data, allowing it to be applied to more cases. Apparent rotation period calculations can be extended to calculate, for example, inertial periods in some cases, further extending their direct application to attitude estimation. In addition, these methods can include the use of laser data, increasing the possibilities for data fusion. As for attitude estimation, some improvements are already underway, starting with its extension to the use of RCS measurements (for which GRIAL is currently being enhanced). Although as a second priority, the use of moving parts in the simulator is also envisaged in short term, as well as to include the use of laser measurements. Going more into the details of the method, additional improvements would be using different filters (e.g. Kalman), optimising the parameterisation of the algorithm and considering uncertainties in the a-priori information to analyse their impact on the attitude estimation.

However, being a complex problem, not definitively solved at present, the presented approach can be considered as a promising starting point as a methodology for object characterisation.

ACKNOWLEDGEMENTS

We would like to thank the European Commission and the SAURON Project Consortium, as many of the algorithms within this methodology have been developed in the frame of this project. This project has received funding from the European Defence Industrial Development Programme (EDIDP) under grant agreement No EDIDP-SSAEW-SSAS-2020-086-SAURON. The work performed reflects only the authors' view and the Commission is not responsible for any use that may be made of the information it contains.
We would like to thank the GMV team that developed GRIAL, to allow us to use it and support our developments within this study.
We would like to thank Mini-MegaTORTORA (MMT9). Mini-MegaTORTORA belongs to Kazan Federal University.

REFERENCES

[1] Jiří Šilha, Matej Zigo, Tomáš Hrobár, Peter Jevčák, Martina Verešvárska. Light curves application to space debris characterization and classification, *8th European Conference on Space Debris (virtual)*, April 2021

[2] Matthew A. Stevenson, Michael Nicolls, Chris Rosner. Space Object Attitude Stability Determined from Radar Cross-Section Statistics, *AMOS*, 2019

[3] Roberto Furfaro, Richard Linares, Vishnu Reddy. Space Objects Classification and Characterization via Deep Learning and Light Curves: Applications to Space Traffic Management, *Space Traffic Management Conference*, 2019

[4] Emma Kerr, Elisabeth Geistere Petersen, Patrick Talon, David Petit, Chris Dorn, Stuart Eves, Using AI to Analyse Light Curves for GEO Object Characterisation, *AMOS* 2021

[5] The Consultative Committee for Space Data Systems, CCSDS Recommended standard for Tracking Data Message, CCSDS 503.0-B-2, Blue Book June 2020 - URL https://public.ccsds.org/Pubs/503x0b2c1.pdf

[6] European Space Agency, Database and Information System Characterising Objects in Space (DISCOS), [Online; accessed 25-Aug-2023] (2023) - URL https://discosweb.esoc.esa.int/

[7] T. Kelso, CelesTrak, [Online; last accessed 25-Aug-2023] (2023) - URL https://celestrak.org

[8] 18th Space Control Squadron, Space-Track, [Online; last accessed 25-Aug-2023] (2023)
URL https://www.space-track.org

[9] C. J. Wetterer, M. Jah, Attitude estimation from Light Curves, *Journal of Guidance, Control and Dynamics, Vol-32 N.5*, 2009

[10] R. Linares, J. L. Crassidis, M. Jah., Particle filtering light curve based attitude estimation for non-resolved space objects, *Advances in the Astronautical Sciences* 152:119-130, 2014

[11] R. Linares, J. Crassidis, M. Jah, H. Kim, Astrometric and photometric data fusion for resident space objects orbit, attitude and shape determination via multiple model adaptive estimation, *AIAA Guidance, Navigation and Control Conference*, 2010

[12] Paulete, Carlos, et al. AIMLRCS: A Machine Learning approach to spacecraft attitude and object identification based on RCS from the S3TSR. *8th European Conference on Space Debris*. 2021. p. 167.

[13] Hochreiter, Sepp, and Jürgen Schmidhuber. "Long short-term memory." Neural computation 9.8, 1997

[14] Karpov et al, "Massive photometry of low-altitude artificial satellites on Mini-Mega-TORTORA", *Revista Mexicana de Astronomía y Astrofísica (Serie de Conferencias)* Vol. 48, pp. 112-113 2016

[15] G. Quint, A. de Andrés, M. Viturro, J. Carro, V. Morand, M. Steindorfer et al. An Advanced Tool to Determine the Apparent Rotation Period of a Space Object from a Fusion of Measurements, *IAC-22-68573*, 2022.

[16] Silha, J., et al. Apparent rotation properties of space debris extracted from photometric measurements, 2018.

[17] Balachandran, K. et Subbarao, K. Estimating Sidereal rotation period of Resident Space Objects using non-uniformly sampled light curves, *AMOS*, 2018

[18] D. Kucharski et al., Attitude and spin period of space debris Envisat measured by satellite ranging, *IEEE Transactions on Geoscience and Remote Sensing*, 2014.

[19] G. Kirchner, W. Hausleitner, and E. Cristea, Ajisai spin parameter determination using Graz kilohertz satellite laser ranging data, *IEEE Transactions on Geoscience and Remote Sensing*, 2007

[20] VanderPlas, Jacob T. Understanding the Lomb-Scargle Periodogram, *The Astrophysical Journal Supplement Series*, 2018.

[21] S. Larsson. Parameter estimation in epoch folding analysis, *Astronomy & Astrophysics Supplement Series 117, 197-201*, 1996.

[22] Cognion, Rita L. Rotation rates of inactive satellites near geosynchronous earth orbit, *AMOS*, 2014.

[23] Gao, F. and Han, L. Implementing the Nelder-Mead simplex algorithm with adaptive parameters, *Computational Optimization and Applications*, 2012.

[24] Hejduk, M.D., Specular and Diffuse Components in Spherical Satellite Photometric Modeling, *AMOS*, 2011.

[25] Y.-l. Xu and C. Stokely. A Statistical Size Estimation Model for Haystack and HAX Radar Detections, *56th International Astronautical Congress*, 2005.

[26] M. K. Mulrooney, M. J. Matney, M. D. Hejduk, E. S. Barker, An investigation of global albedo values, *AMOS*, 2008.

[27] Murray, J., Haystack Ultra-Wideband Satellite Imaging Radar Measurements of the Orbital Debris Environment, 2020.

[28] National Aeronautics and Space Administration, LAGEOS: LAser GEOdynamic Satellite [Online; last accessed 25-Aug-2023] (2023) – URL https://lageos.gsfc.nasa.gov/

[29] M. Pearlman, D. Arnold, M. Davis, Laser geodetic satellites: a high-accuracy scientific tool, *Journal of Geodesy*, 2019.

[30] B. D. Tapley, B. E. Schutz, G. H. Born, Statistical Orbit Determination, *Elsevier Academic Press*, 2004

[31] M. Hejduk, H. Cowardin, E. G. Stansbery, Satellite Material Type and Phase Function Determination in Support of Orbital Debris Size Estimation, *NASA Johnson Space Center*, 2012.

[32] Cook, R. L., & Torrance, K. E. (1982). A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, *1*(1), 7-24.

[33] Walter, B., Marschner, S. R., Li, H., & Torrance, K. E. (2007, June). Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (pp. 195-206).

[34] Fairbairn, Maxwell B. "Planetary photometry: The lommel-seeliger law." *Journal of the Royal Astronomical Society of Canada, Vol. 99, No. 3, p. 92* 99 (2005): 92.

[35] Khronos Group, OpenGL, [Online; last accessed 28-Aug-2023] (2023).

[36] Le calcul de la transmission spectrale atmosphérique, [Online; last accessed 28-Aug-2023] (2023)
URL http://www.astrosurf.com/buil/extinction/calcul.htm

[37] Mendes, V. B., and Erricos C. Pavlis. "High-accuracy zenith delay prediction at optical wavelengths." *Geophysical research letters* 31.14 (2004).

[38] European Space Agency, SMOS [Online; last accessed 28-Aug-2023] (2023)
URL https://earth.esa.int/eogateway/missions/smos

[39] European Space Agency, Sentinel-3 [Online; last accessed 28-Aug-2023] (2023)
URL https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-3

[40] Gerber, J., et al. "Sprint Advanced Concept Training (SACT): A renaissance in collaborative international space operations." *AMOS*. 2020.