

ARTIFICIAL INTELLIGENCE FOR SPACE RESIDENT OBJECTS CHARACTERISATION WITH LIGHTCURVES

Keiran McNALLY⁽¹⁾, Diego RAMIREZ⁽¹⁾, Alfredo M. ANTON⁽¹⁾, Duncan SMITH⁽²⁾, James DICK⁽³⁾

⁽¹⁾GMV NSL Ltd. Harwell, HQ Building, Thomson Avenue Harwell Campus Didcot OX11 0GD, UK, kmcnally@gmvnsl.com, dramirez@gmvnsl.com, amanton@gmvnsl.com

⁽²⁾Beechleaf Consulting. The Old Exchange, Church Street, Castle Cary, Somerset, BA7 7EJ, UK, duncan.smith@beechleafconsulting.com

⁽³⁾Space Insight Ltd. PO Box 2993, Eastbourne, BN21 9BD, UK, jsbd@spaceinsight.co.uk

ABSTRACT

The power of machine learning (ML) on classification problems has been well documented, from image classification to speech recognition, so its use in lightcurves seems reasonable as a new field of applicability. Lightcurves can be used to identify not only the attitude state, but also the type of platform, the orientation, and geometry. The space domain presents a challenge to the applicability of ML due to the large amount of training data required. The creation of a lightcurve simulator utilising the well-known Phong reflection model allows for the simulation of thousands of well-characterised training samples. Initial tests show that a neural network can identify the satellite and attitude mode from lightcurves with an accuracy of 95% using simulated data. Testing on real data shows care must be taken to ensure that the simulated data includes realistic gaps and noise; including these features, leads to an accuracy of at least 70%.

1 LITERATURE REVIEW

Utilising existing photographic and astronomical equipment for artificial satellite observations began with the advent of the satellite era in the late 1950s and has been evolving quickly since. Starting from photographic film, the state of the art moved to photomultiplier tubes [1], then to various generations of charge couple devices (CCD) and now to extending observational capabilities beyond the visible light spectrum.

Optical tracking data, obtained from either photometers or non-resolved images, has historically been used for the orbit determination of resident space objects (RSOs). Recent research has proven that a collection of photometric measurements, i.e. the record of object brightness in time (traditionally referred to as a lightcurve) can provide vastly more detailed information, making it possible to infer the spin and attitude dynamics, shape and, to some extent, surface parameters.

1.1 Spin and Attitude State Characterisation

Analysing lightcurves for intrinsic periodicities and their changes is the most fundamental approach to

characterising space objects via photometric observations. Since it is a general problem of time series analysis, several methods and techniques have been developed and commonly adopted in the space situational awareness (SSA) field.

Fourier Transforms, in particular the Fast Fourier Transform (FFT), are commonly used to determine the spectral density when searching for dominating periodicities of time series data. The main drawback of any Fourier Transform (FT) method is the requirement of an equally spaced time series that is free of gaps and strong noise, conditions that are rarely present in satellite observations. Various pre-processing techniques have been applied to the input data to mitigate mentioned shortcomings, however FT-based methods are commonly used as a crosscheck for other periodicity searching techniques.

On the other hand, a method specifically designed for sparse and unevenly sampled data is least squares spectral analysis, Reference [2] observes that rotation periods up to 40% of the duration of the lightcurve can be detected using the *Lomb-Scargle* algorithm only if the data is free of gaps. Introducing gaps into the data reduces the duration of the detectable rotational period to just 16% of the total time series. The same reference also shows that *Epoch Folding* is able to detect periods if the length is less than 45% of the total lightcurve.

Typically, a combination of several methods, in addition to the individual experience is used to extract reliable spin and attitude state information from lightcurve data.

1.2 Shape Estimation

The Light Curve Inversion (LCI) approach aims to determine approximately the shape of a satellite and is one of the key aspects of ongoing space situational awareness research efforts. Using photometric data with a light curve inversion technique commonly used for asteroid shape estimation; [3] presents an initial investigation of the approximate shape of man-made objects in the GEO belt. The results indicate that light curve inversion is a beneficial technique for rocket bodies and CubeSats, while presenting only moderate

difficulties with high area-to-mass ratio objects and box-wing shapes.

Reference [4] applies deep learning for RSO classification using photometric data, in preparation for future RSO catalogues, which will have more restraint requirements and shall provide a detailed picture of characteristics such as shape, attitude, angular velocity, as well as surface parameters. Noting that physically based model inversion (via Extended Kalman Filters) is computationally expensive, the authors show how deep learning methods (Convolutional Neural Networks) can be deployed to provide an effective shape retrieval in a fast and accurate fashion. Using purely simulated data for four simple models representative of typical rocket body shapes, shown in Fig. 1-1., an average classification accuracy of over 90% is achieved, with a maximum confusion of 2% between classes “a” and “d”. This work serves as a proof of concept regarding the use of machine learning for RSO classification.

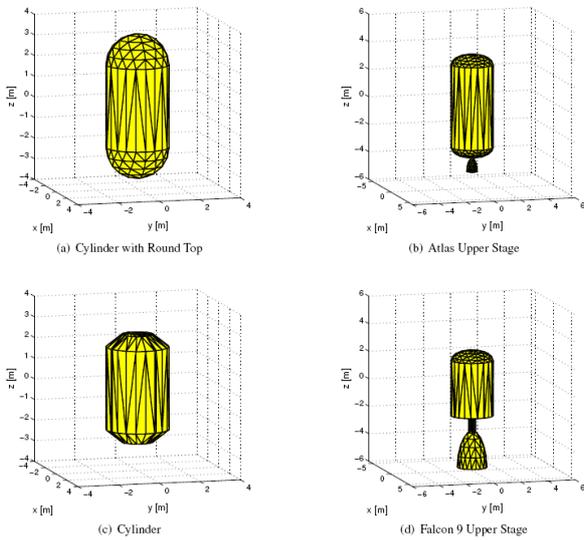


Figure 1-1. Simulated rocket body models as seen in [4].

1.3 Phong Reflection Model

The mathematical model chosen for use in the light curve simulator was the Phong reflection Model, which has been widely used in computer graphics since it was developed by Bui Toung Phong in 1975. This model belongs to a larger collection know as bidirectional reflectance distribution functions (BRDFs) which are used to calculate the brightness of a point based on the directions of the light source and the observer.

The brightness computed from the Phong model is the result of three components: ambient, diffuse, and specular reflections.

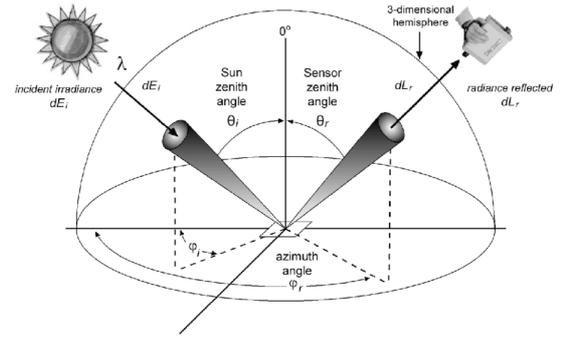


Figure 1-2. Breakdown of a BRDF. [5].

The ambient brightness is the uniform light distribution within a scene, the diffuse component is the light reflected in all directions by a point, and the specular component consists of bright highlights due to light reflecting off a shiny surface. A simplified visualisation of these components is presented in Fig. 1-3.

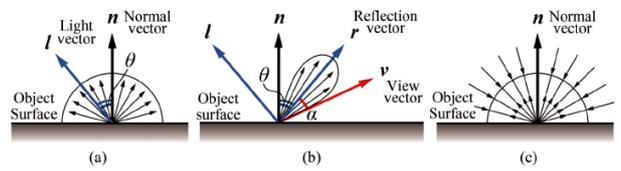


Figure 1-3. Phong Reflection Model. (a) diffuse reflection light, (b) specular reflection light, and (c) ambient reflection light [6].

For satellites in Earth orbit, the problem can be simplified by assuming that the ambient component is negligible and the only light source present in the scene is the Sun.

This reduces the reflection model to:

$$I = k_d(\hat{L} \cdot \hat{N})i_d + k_s(\hat{R} \cdot \hat{V})^\alpha i_s \quad (1)$$

Where:

- I , is the reflected intensity,
- $k_{d/s}$, are the material coefficients for diffuse or specular reflection, respectively,
- $i_{d/s}$, are the diffuse and specular intensities of the light sources, respectively,
- α , is a parameter that determines the “shininess” of the material (with shiny, mirror like surfaces having a higher value),
- \hat{R} , is the direction vector of perfect reflection and is defined as:

$$\hat{R} = 2(\hat{L} \cdot \hat{N})\hat{N} - \hat{L} \quad (2)$$

Where:

- \hat{L} , is the direction vector from the point to the light source,
- \hat{N} , is the normal vector at the given surface point,

- \hat{V} , is the direction vector from the point to the viewer.

Due to the separation of the observer from the satellite, it can be assumed that the direction vectors of both the light source and the viewer are constant across the reflective surface, thus the brightness is dependant only on the surface normal (since \hat{R} is only dependant on \hat{N}). It is therefore relatively simple to integrate on each surface and obtain the brightness of a 3-D shape.

It can be observed that the diffuse component is independent of the viewer direction and is determined by the angle between the surface normal and the light source. For a satellite, this means that the diffuse brightness will vary slowly with the position and attitude of the satellite with respect to the Sun. The specular component, on the contrary, depends on all the position vectors and is responsible for the characteristic peaks that are often seen in lightcurves.

The ambient component of an object's appearance is the result of scattered (and often multiply scattered) light, which comes from an object's surroundings rather than directly from a light source [7]. Ambient light in the context of a room on Earth would be light scattered by the structure of the room (e.g. the walls and ceilings) and other objects within the room. For objects in space, the ambient light would be from distant stars, galaxies, nebulae, etc. (with the exception of earthshine, which is discussed separately).

The radiance of cosmic illumination has been measured on the surface of the Earth [8]. The value can be used as an approximation of the value that would be seen in Earth orbit. The approximate value is quoted by [8] as $\sim 2.2 \cdot 10^{-4}$ lux as the radiance on a flat surface due to the cosmos. In comparison, the Sun would contribute $\sim 10^{+5}$ lux, and the full Moon $\sim 10^{-1}$ lux. While the values are approximations, the radiances from the cosmos and the Moon are orders of magnitude less than the direct solar illumination and so will not contribute significantly to a satellite's photometric signature.

Sunlight reflected from the Earth (earthshine) is often visible as a brightening of the full disk of a thin-crescent Moon. This phenomenon is most prominent in the evening, when the Moon is a few days past new and sets not long after the Sun. The amount of earthshine from an object depends on how much of the sunlit Earth is visible to the surfaces of the object, which are visible to the observer. The extent of planetary cloud cover and the scattering characteristics of the sunlit land and cloud will all be significant – but difficult to model – factors, which affect the amount of earthshine, directed towards the object's visible surfaces.

GEO satellites are typically observed at, or near, opposition (see Fig 1-4) where the object is at considerable angular distance from the Sun, the Earth-

facing side of the object will be facing the night side of the Earth and so no Earth-reflected reflected sunlight will illuminate the object. Modelling indicates earthshine will be less than 10^{-3} that of the direct solar illumination [9].

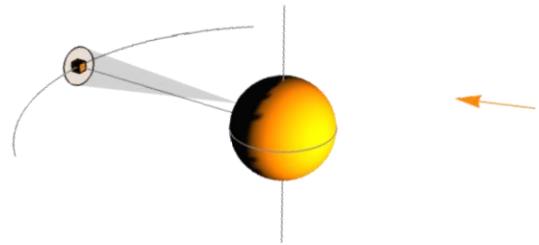


Figure 1-4. Satellite viewed at opposition (arrow indicates direction of sunlight).

2 LIGHTCURVE SIMULATOR

Due to the lack of large amounts of well-characterised real photometry data to use in training machine learning algorithms, it was decided in this project to implement a simulator capable of producing large amounts of realistic lightcurves. The simulated lightcurves were validated against examples from the archive of SpaceInsight Ltd (SIL) and the expertise of Beechleaf Consulting Ltd (BCH) was used to recreate, as accurately as possible, the model of the satellites used in validation.

An extensive collection of satellite images sourced by BCH, were used to provide the configuration of each satellite, including the number of antennae and the approximate size of each shape. The material properties of each surface were also estimated using these images, combined with baseline values for typical materials provided by both SIL and BCH.

The satellites chosen for final validation were a result of iterative work comparing the satellites available in the SIL archive and those that could be easily modelled by use of artist impressions and photographs.

The lightcurve simulator was designed with the goal of producing sufficient data with which to train a machine-learning algorithm. Computationally intensive methods such as 3D modelling and ray tracing were discarded, and finally the Phong reflection model (presented in Section 1.3) applied to satellite models composed of simple shapes was considered reliable and accurate enough for the purpose. The simple shapes composing those models consisted of cubes, flat plates, cylinders, and spheres. Each shape is treated individually and is considered to be at the origin of the satellite model and so any effects from shadowing or multi surface reflections are not considered. Each shape is created by specifying the size and material properties: ambient coefficient, diffuse coefficient, specular coefficient and specular index. The satellite can be configured as either spinning or stable

with configurable pointing offsets or spin rates and each shape can be offset from the main satellite frame to recreate antenna pointing.

The simulator also includes the capability to randomise several parameters per lightcurve in order to recreate variations to the configuration and the location of the spacecraft. The key parameters that can be randomised within a configurable margin are:

- Satellite longitude
- Satellite body pointing
- Shape fixed orientation (w.r.t. the satellite body).
- Shape initial orientation and spin rate.
- Material coefficients
- Specular index

2.1 Validation with SIL Archive

In order to verify the satellite models created for the simulation, full night photometry data provided by SIL was used. It was decided that this study would focus on GEO satellites (due to the full night photometry available) and so this archive is well suited as it contains mostly GEO photometry recorded by the Starbrook optical telescopes located in Troodos, Cyprus. This data was for satellites identified by BCH as well defined and of similar box wing designs. The photometry files were all recorded in September 2018 and due to the alignment of the Sun and Earth at this time of year the shadow of the Earth was present, which caused a gap of approximately 70 mins and so some definition was lost. Comparisons between the simulated and real photometry are shown in Figs 2-2 and 2-3.

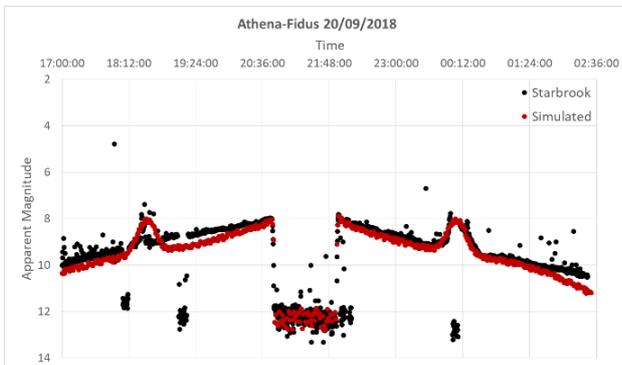


Figure 2-1. Athena-Fidus Starbrook Comparison

The data gap that is seen at roughly the midpoint of the previous lightcurves was caused by the Earth's shadow, and so there is no illumination from the Sun. The values shown are the noise caused by the background subtraction process. This functionality was included in the simulator and as such, the background noise can be configured.

The smaller areas of noise present in the Starbrook

photometry are mainly due to clouds or other obscuring phenomena and so cannot be recreated accurately in the simulator due to their unpredictability.

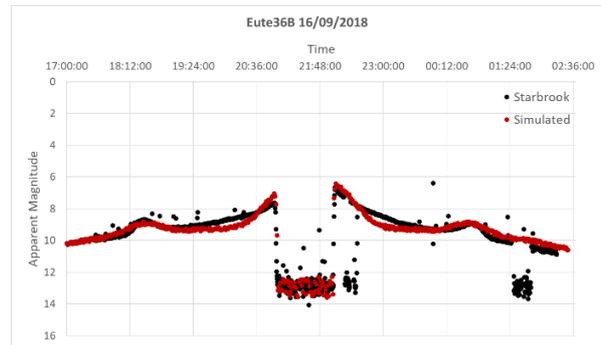


Figure 2-2. Eutelsat 36B Starbrook Comparison

Towards the end of the lightcurve in Fig. 2-2, the apparent magnitude of the simulated data can be seen to drop off compared to the observed photometry. The cause of this is unclear. It is possible that the simple shapes used to create the satellite were missing details only observed around twilight, or that the photometry reduction methods used do not cope well with effects seen around twilight; other factors such as the presence of the galactic plane in the image could also contribute to errors seen. Nevertheless, it has been demonstrated that the simulator is able to provide representative lightcurves using a relatively simple reflection model and basic shapes.

Tumbling satellites are more difficult to model due to the uncertainty regarding the spin axis, spin rate and the deployment of panels before the tumbling state began. Photometry of tumbling satellites is far scarcer than for stable satellites and as such only two tumbling objects were initially available from the SIL archive - Angosat 1 and GSat 6A. For each satellite, only one photometry file was available, in comparison to an average of 10 for stable satellites. The experience of SIL was used to estimate the configuration of each tumbling satellite including the spin rates and satellite configurations. The comparison between the real photometry (black lines), SIL simulated photometry (blue lines), and the lightcurve simulator (red lines) is shown in the following Fig. 2-4.

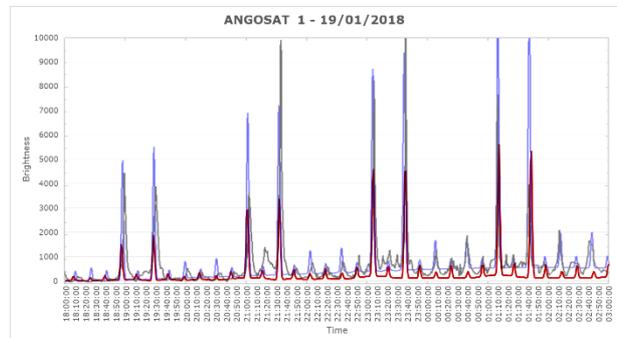


Figure 2-3. Angosat 1 Photometry Comparison

It is possible to obtain a realistic approximation of the spin rate by examining the real photometry and making some assumptions. For Angosat 1 it was assumed that the twin peaks seen in Fig. 2-4 were caused by the two side mounted antenna and thus the rotational period can be estimated as approximately two hours. It can be seen in Fig. 2-4 that the frequency of the peaks is well matched by both simulated lightcurves, as is the relative size of each peak. The absolute value of each peak is not recreated well however with both simulated lightcurves either under or over-estimating different peaks. It is therefore likely that a more detailed simulation method is needed to capture effects that are not currently considered, such as overlapping shadow from multiple surfaces. The machine learning testing with real data will be solely focussed on real data due to the difficulty in retrieving multiple lightcurves for any given tumbling satellite and so the current simulation method is well suited for this purpose.

3 AI APPROACH

Based on the available expertise and the literature review, three distinct approaches were devised, each with an alternate approach to processing the data prior to the neural network.

- The first approach involves doing no processing to the raw data (brightness), only normalisation when needed, and simply using the raw lightcurve as the input. This will allow the neural network to decide what features are linked to changes in brightness and how the phase angle can influence the brightness.

- The second approach involves deciding what features are likely to influence the lightcurve and only provide this to the AI, greatly reducing the volume of data required. Data processing for this method involves applying pre-processing methods to the raw lightcurves, such as polynomial fitting, providing the coefficients and residual as input, or the frequencies identified by a Fourier transform.

- The third approach is taken directly from the literature [10], where it was applied to sparse photometry of distant solar bodies by taking the difference in brightness and time of every combination of pair of points. This idea was presented as a way to extract patterns of sparse, uneven data but is applied here to simulated lightcurves.

3.1 Interfaces and pre-processing

To perform any needed data handling and normalisation, an interface was created to process the simulated data before it was passed as input into the neural network. The main function of this interface is to collate all the lightcurves into a single csv file, which can be easily loaded and then reshaped to form the inputs to the neural network. Prior to creating this csv file, the interface applies any normalisation, filtering, pre-processing

and/or other configurable data augmentation techniques. In addition, this interface standardises the real data to the size required by the input layer of the neural network, and removes possible outliers.

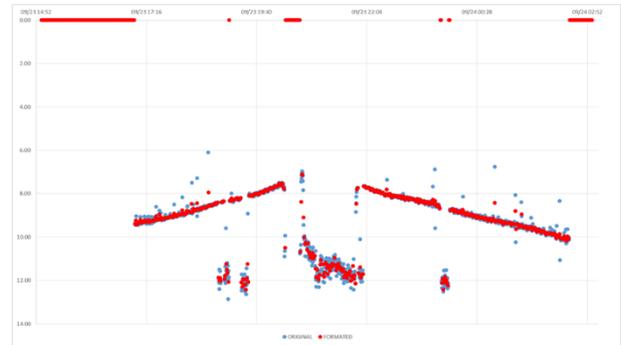


Figure 3-1. Original and formatted apparent magnitude comparison (2003-020A)

3.2 Raw Data

This approach is the simplest of the three considered. It involves passing the lightcurve from the simulator into the machine learning algorithm with no additional data processing or manipulations applied, apart from the ones related to the data interfaces.

The data was normalised using two methods: 1) to have zero mean and unit variance and 2) re-arranged to ensure all the outputs lie between zero and one.

$$x_{normalised} = \frac{x - mean(x)}{std\ dev(x)} \quad (3)$$

For testing with real data, the same normalisation was applied but using the mean and standard deviation from the training data to ensure the same distribution and range of values is considered. A combination of both normalisations were used and it was found that the best results when testing with real data were achieved when both normalisation methods were applied. When dealing with simulated data only, the best results occurred with either both normalisation methods or none applied. The use of only one of the normalisation methods not only yielded poorer accuracy but also took more training epochs to converge to a stable accuracy.

3.3 Feature Analysis

This approach involves the analysis of the lightcurves looking for the most important characteristics or features. In this case, only a preliminary analysis was performed, including maximum and minimum values, the possible polynomial coefficients of a polynomial fitting and the most representative frequencies of its Fourier transform. During the analysis, no real data was used, as the other methods discussed in this section displayed more promising results.

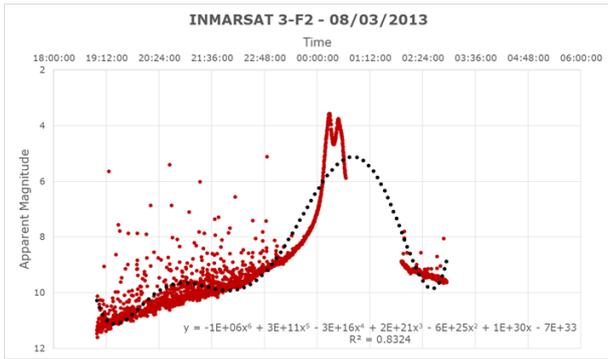


Figure 3-2. Schematic representation of the polynomial fitting

3.4 Differential Approach

This approach was taken from [10] and it is designed to capture the variability in the lightcurve by showing the variation in brightness and time, instead of the absolute values. The computation involved taking the first value of brightness and time and computing the difference in each for every other point in the lightcurve. This was then repeated for every value of apparent magnitude and time leading to $(x - 1)^2$ pairs of points, where x is the number of points in the original lightcurve.

This method was originally applied to photometry taken from the Catalina Real-Time Transient Survey, which contained optical measurements of transient celestial phenomena such as supernovae. The differential approach was applied here due to the uneven and sparse nature of the data, with the idea being to create more data while capturing the variability of the original.

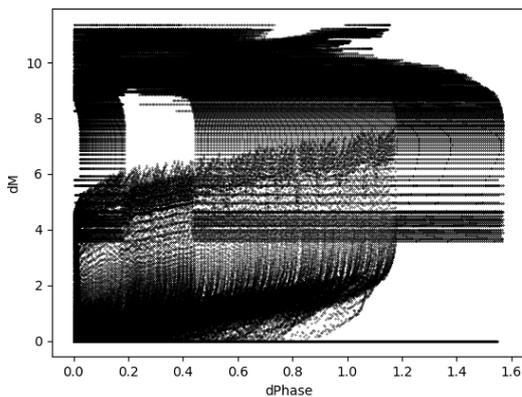


Figure 3-3. Example of an input for the differential approach

This differential mapping was applied in the same way to the simulated lightcurves used for the initial machine learning testing where the time component was replaced by the phase angle. This generated 1,036,800 data points per lightcurve, with half being the brightness difference

and the other half being the corresponding phase angle difference. Given that machine learning datasets could contain thousands or millions of training examples, it was decided to sub-sample (1/10) the simulated datasets to perform the initial tests.

4 RESULTS AND DISCUSSION

In order to understand the precision in the predictions of each algorithm, the confusion matrix tool was used. This graphical representation allows discerning how far the AI predictions (the ordinates' labels) are from reality (the abscissae's labels). The darker the diagonal the better the prediction.

For each of the approaches discussed here, the same satellites, shown in Tab. 4-1., were used as a dataset unless otherwise stated.

Table 4-1. Satellites Used

Sat. Name	Sat No.	Attitude Mode
Athena-Fidus	39509	Stable
Eutelsat 36B	36101	Stable
Hellas-Sat 2	27811	Stable
Sicral 2	40614	Stable
Angosat 1	43087	Tumbling
GSat 6A	43241	Tumbling

4.1 Raw Data

4.1.1 Simulated Data

Initial testing was carried out using only simulated data and using only apparent magnitude as the input. The neural network used to obtain the best results for this phase is shown in Tab. 4-2. Note that a feature here is defined as a unique input variable to the neural network. For example, a lightcurve consisting of apparent magnitude and phase angle at every time step would have two features at every epoch.

Table 4-2. Baseline Raw Data Neural Network Design

Layer	No. of Neurons
Input	721 x no. of features
Dense fully connected	2000 x no. of features
Dense fully connected	2000 x no. of features
Dense fully connected	2000 x no. of features
Flatten	2000 x no. of features
Output	no. of classes

Using this design and training the network for 100 epochs (iterations), the neural network achieves 88% accuracy on identifying the satellite using apparent magnitude only.

Fig. 4-1 shows the confusion matrix for this scenario, which reveals a high level of confusion between the stable satellites, in particular Eutelsat 36B and Hellas-Sat 2, with up to 19.2% and 24.4% confusion in labelling each class. Examining the artist's impressions shown in Fig. 4.2 shows that these two satellites have a similar configuration of front mounted antenna, which may be causing this significant confusion. Although Eutelsat 36B appears to have four side mounted antenna whereas Hellas-Sat 2 has only two, the total surface area may be similar and thus adding to the similarities. Athena-Fidus appears to have the most distinctive profile with six small antenna front mounted antenna and three side mounted. The result of this is clearly seen in Fig. 4-1 as Athena-Fidus has the highest classification accuracy of all the stable satellites.

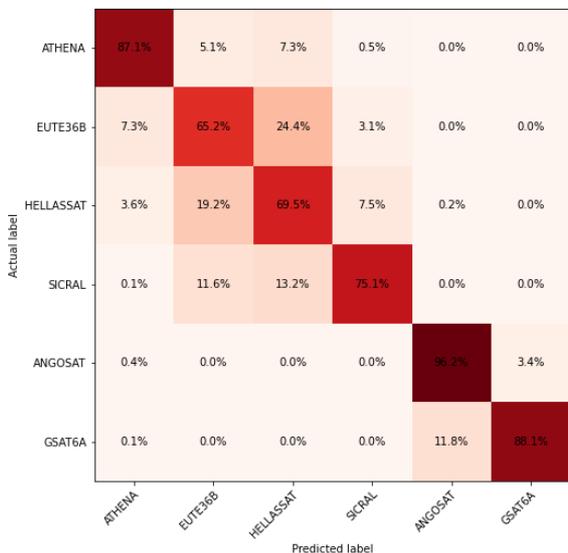


Figure 4-1. Satellite Identification Confusion Matrix (Apparent Magnitude Only)



Figure 4-2. Stable Satellites

The neural network also shows an accuracy of 88-96% in distinguishing Angosat 1 from GSat 6A, which are the two tumbling satellites in this dataset. Both of these satellites are tumbling because of a failure during GEO orbital insertion. The spin rates and initial positions for each are randomised within the same margin and so the only difference between the satellites is the construction of the models. Fig. 4-3 shows the artist impressions for Angosat 1 and GSat 6; GSat 6A was of identical design but suffered a failure during orbital insertion and the large antenna was never deployed. As such, there is no antenna included in the GSat 6A model whereas Angosat 1 is modelled as having two side-mounted antennae, which may account for the relatively high accuracy in identifying between these two satellites.

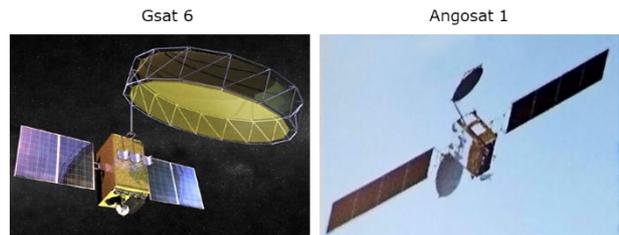


Figure 4-3. Tumbling Satellites

Including the phase angle as part of the input layer increases the accuracy of the network to 95%, as is shown in Fig. 4-4 where the confusion matrix is more strongly aligned with the diagonal. Including the phase angle may be allowing the neural network to identify patterns in how the brightness is dependent on the phase angle and so any anomalies caused by surface offsets may be easier to detect, hence the increase in accuracy.

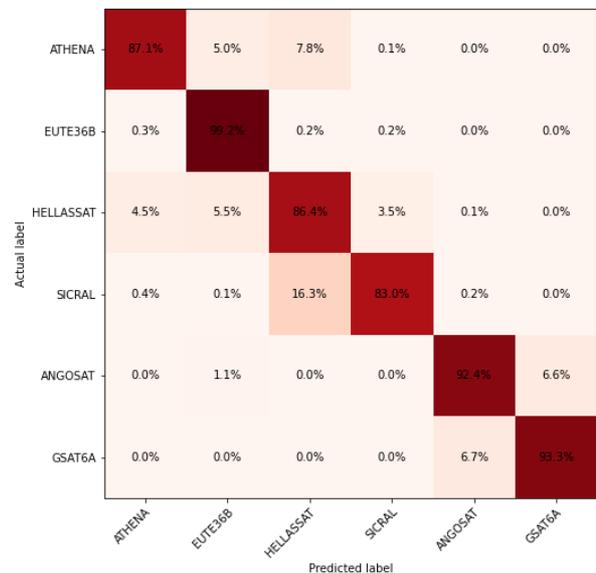


Figure 4-4. Satellite Identification Confusion Matrix (Apparent Magnitude + Phase Angle)

The results from Figs 4-1 and 4.4 show very little

confusion between the tumbling (Angosat 1 and GSat 6A) and stable satellites with a maximum confusion of only 1%. This was analysed in more detail by taking one stable satellite (Athena-Fidus) and one tumbling satellite (Angosat 1) and simulating each as both stable and tumbling. The results show an accuracy of over 95%, which would suggest that the neural network is clearly able to identify which lightcurves contain periodicities and which do not.

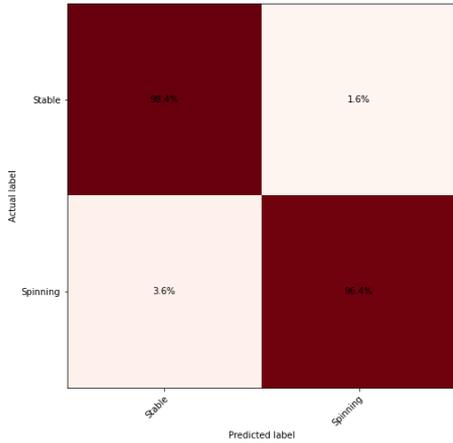


Figure 4-5. Attitude Mode Confusion Matrix

Preliminary work was conducted to test the ability of the neural network to identify the current pointing of satellite antenna. This could have real world applications such as fault identification or aiding in the classification of unknown satellites by determining where the antenna are pointing on the Earth. For this, a simple satellite was created using Hellas-Sat 2 as a basis. This mock up satellite was only modelled with one large antenna of 5m² (compared to other modelled antenna of 0.5-2m²). This antenna was offset in pitch (North - South) and then yaw (East – West) in increments of 0.1 radian (5.73 degrees) from -0.4 to +0.4 radians. It was found that the neural network was unable to recognise the pitch variations but showed remarkable accuracy in identifying the yaw offset in certain circumstances. Two scenarios are presented here: the first consists of lightcurves that have been simulated on a random day of the year and the second where the lightcurves are constrained to be on a random day within one month. The unconstrained scenario yields an accuracy of 30-40% and considerable confusion between small offset values as shown in Fig. 4-6 although the accuracy increases with a larger offset value, as should be expected. The labels correspond to the specific offset value, where “NEG_POINT_FOUR” is a yaw offset of -0.4 radians. The interesting point about this is the constant confusion with the undisturbed value (ZERO_YAW). Regardless of the actual offset, there is at least 20% confusion with this zero value, although it does appear to reduce slightly as the offset magnitude is increased.

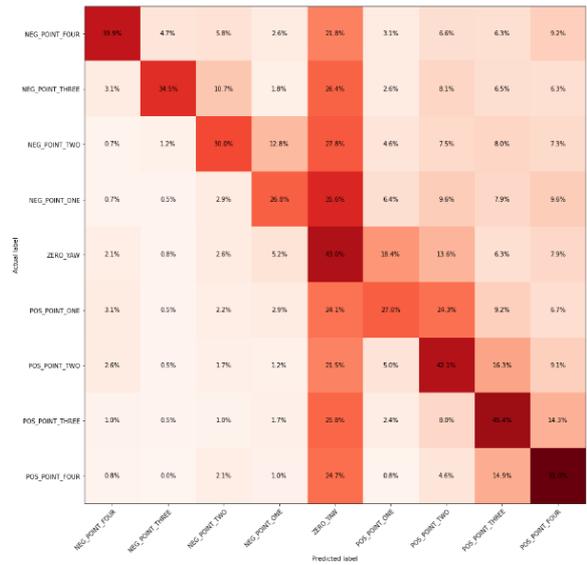


Figure 4-6. Surface Offset Detection Confusion Matrix (One Year)

Given the previous results, a second trial was attempted but limiting the lightcurves to be generated within one month, largely reducing the variability in the lightcurves by limiting the position of the Sun. This allowed an accuracy of 80% to be achieved and removing the previous confusion with the undisturbed value. Fig. 4-7, shows impressive results with over 90% accuracy achieved in identifying the larger offsets. With the offsets close to zero, noticeable confusion can be seen with the nearest labels, mainly in the ±0.5 radian categories but this is quickly reduced as the offsets are increased in magnitude.

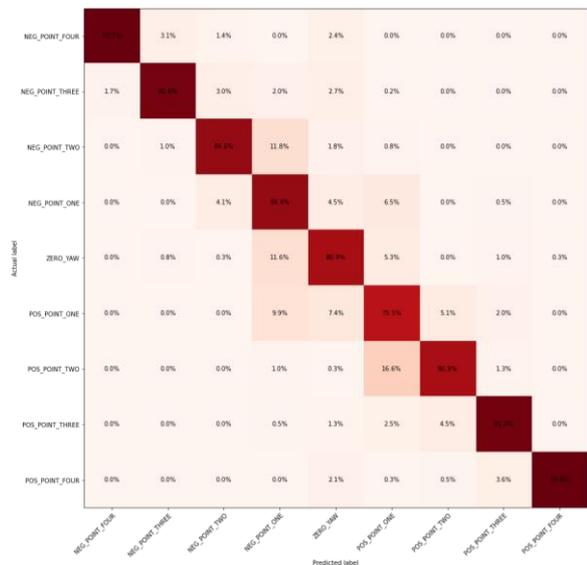


Figure 4-7. Surface Offset Detection Confusion Matrix (One Month)

These results suggest that the effect of the solar declination somehow masks the effects that offsetting one panel can have. This is likely the cause of the confusion seen in Fig. 4-6, which was greatly reduced by limiting the lightcurves to one month. Further work is warranted to explore this effect and to attempt improvements that can somehow decouple the effects of antenna offsets and the solar declination. It is worth noting that both scenarios used a dataset of 2,000 lightcurves and it is likely that increasing this would improve the ability of the neural network to distinguish the antenna pointing effects from the phase angle effects.

4.1.2 Real data

Real data testing was limited to stable satellites due to the lack of sufficient photometry for tumbling satellites; therefore, the only scenario that could be tested on real data was satellite identification. The stable satellites used were Athena-Fidus, Eutelsat 36B, and Sicral 2 each having 10 photometry files available from the SIL archive, all recorded in September 2018. The real data was processed using the interface, described in Section 3.1 in order to be normalised within the same time span as the simulated data and therefore have the same number of inputs as the simulated data. The neural network was trained and validated using simulated data, with the real data being used as a test dataset. The date of the simulated data was also constrained to be within the same month as the real data. Initial results show poor performance with the neural network failing to produce a clear identification using an initial test dataset containing only Athena-Fidus and Eutelsat 36B.

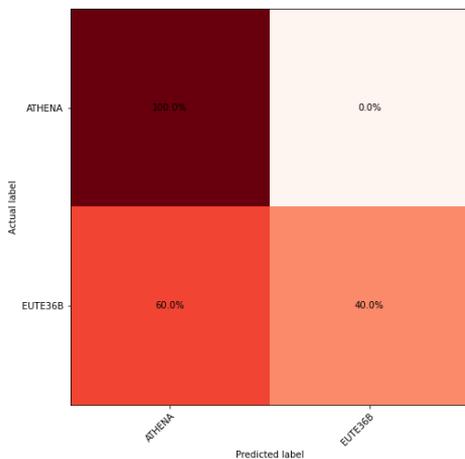


Figure 4-8. Initial Real Data Confusion Matrix

When comparing the simulated and real lightcurves, significant differences were found because of the normalised timespan used. The simulated lightcurves were always produced within a span of local midnight ± 6 hours with a time step of 60s leading to a fixed 721 intervals per lightcurve. The real data can be of any length and with an unknown time step; this is why the

real data interface discussed in Section 3.1 was developed, to normalise the lightcurve length and time step. This caused blocks of zeros to be inserted since the real lightcurves are usually shorter than the standard 12-hour duration. This created a scenario for which the neural network had not been trained for (as the training data contained no zeros) and so the tests on real data were failing.

The interface between the simulated data and the neural network was subsequently updated to insert an array of zeros at the start of each lightcurve to allow training with more representative data. This led to improved results with an accuracy of 70% but the confusion matrix in Fig. 4-8 shows that the majority of examples are classified as Athena-Fidus.

The simulated data interface was then further updated to insert an array of zeros at the start and end of each simulated lightcurve. The length of these zero arrays was randomised with a configurable upper and lower limit. The size of the neural network was increased to account for the loss of significant amounts of training data. Adding dropout to the initial three layers was also found to provide results that were more reliable during validation and testing.

The resulting design is shown in Tab. 4-3. With these updates, the neural network was able to achieve an average of 70% accuracy which varied by $\pm 10\%$ between subsequent executions. This is most likely a consequence of the limited availability of real data.

Table 4-3. Real Data Neural Network Design

Layer	No. of Neurons
Input	721 x no. of features
Dense fully connected	2000 x no. of features
Dropout (0.4)	-
Dense fully connected	2000 x no. of features
Dropout (0.4)	-
Dense fully connected	2000 x no. of features
Dropout (0.4)	-
Dense fully connected	2000 x no. of features
Dense fully connected	2000 x no. of features
Flatten	2000 x no. of features
Output	no. of classes

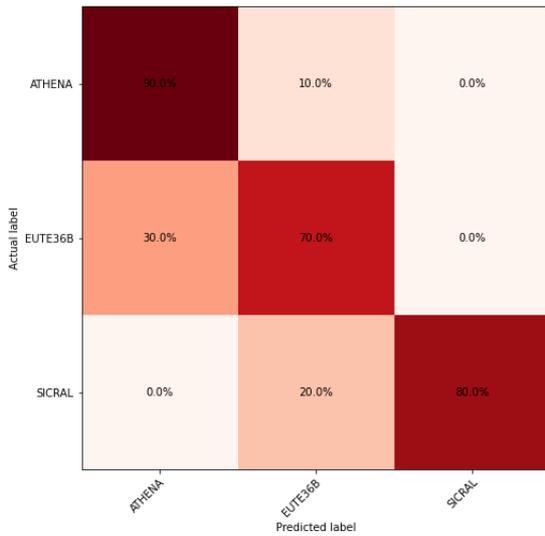


Figure 4-9. Final Data Confusion Matrix

4.2 Feature Analysis

In the case of the feature analysis, the architecture of the neural network used to obtain the best results for this phase is summarised in Tab. 4-4.

Table 4-4: Feature Analysis Satellite Identification NN Design

Layer	No. of neurons
Input	6
Dense fully connected	10
Dense fully connected	8
Flatten	no. of classes
Output	no. of classes

The input layer included the maximum and minimum values of the apparent magnitude, the coefficients of a third order polynomial regression and the first two most representative frequency amplitudes of the Fourier transform.

It is worth to mention the significant differences between the architectures of the previous approach and this approach. Feature analysis used almost 200 times less neurons than the raw data analysis, which influences directly the number of parameters to be trained. This difference is due to the reduced number of inputs used, as, instead of using raw data, the neural network only needs the main characteristics of the lightcurve to identify the satellite.

Using this architecture for the AI and training for slightly less than 100 epochs, the AI was able to achieve almost 80% of accuracy for all the simulated lightcurves as

shown in Fig. 4-10. It can also be seen that tumbling satellites are the most distinguishable satellites of all the ones analysed, reaching more than 99% correct predictions for both satellites. On the other hand, there seems to be a high confusion between stable satellites, leading to correct predictions of even around 30% accuracy. For example, Eutelsat 36B and Sicral 2 only reaches around the 70% of success in the predictions and the AI algorithm could not distinguish correctly between Athena-Fidus from Hellas-Sat 2, as their number of false positives reached a 29.6% and 30.1% respectively.

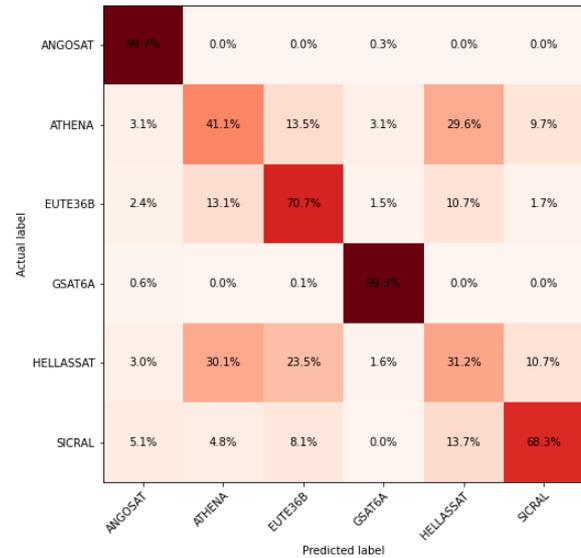


Figure 4-10: Feature Analysis Satellite Identification Confusion Matrix

Fig. 4-3 shows the overall configuration of both tumbling satellites and it is known that GSat 6A had a problem during the launch, which prevented the bigger antenna from being extended. This was taken into account when the simulated data was generated. In addition, both satellites' rotation rates were simulated using different random seeds within a wide range of speeds around their nominal rotation speed. All these facts seem to indicate that this neural network architecture is able to correctly identify tumbling satellites, even with similar spin rates, thanks to the use of the Fourier coefficients, while the polynomial fitting is not covering correctly the behaviour of the stable satellites.

4.3 Differential Approach

Using the differential approach as described in section 3.4, a neural network with the same design as Tab. 4-2 was trained to identify the satellite that created each lightcurve using purely simulated data. The accuracy is comparable to that achieved using raw data (around 90%), but the volume of data used to achieve this accuracy is over 4 times that of the raw data approach, even with the under sampling proposed.

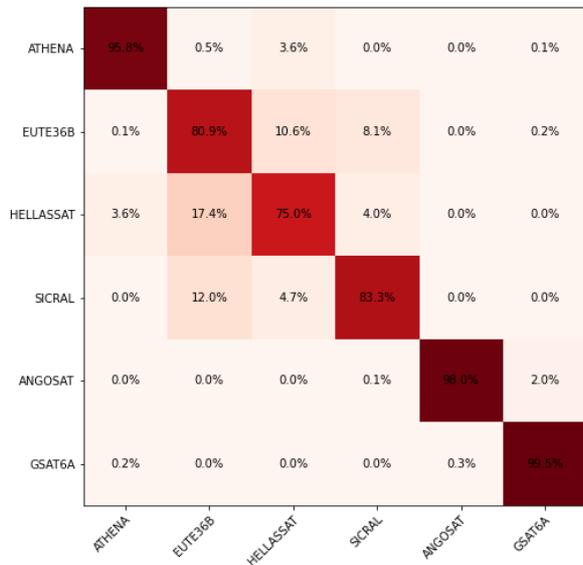


Figure 4-11. Differential Approach Confusion Matrix

The training results for this method are much noisier than the raw approach, with the validation loss varying significantly as shown in Fig. 4.12. There also appears to be a certain degree of overfitting as the validation loss seems to be slightly tending upwards during the latter part of training.

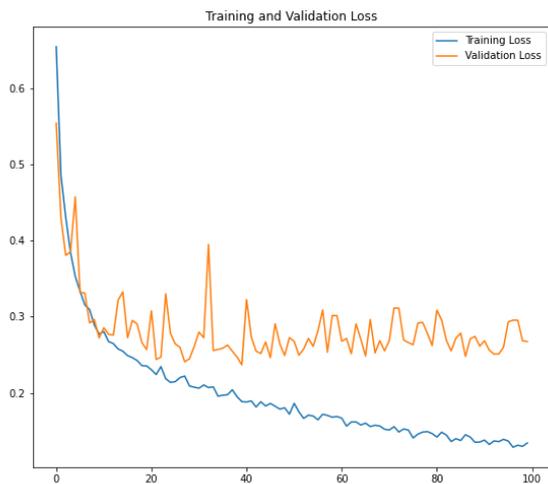


Figure 4-12. Differential Approach Validation Results

5 CONCLUSIONS

The lightcurve simulator solves the problem of the lack of large amounts of real photometry data to use in training a machine-learning model, as it generates photometry data similar to the real data with enough variability to consider a huge number of conditions.

The three approaches considered achieved reasonable accuracy, in particular the feature analysis performed best on the classification of tumbling satellites, while the

other two methods reached at least 90% accuracy in identifying the satellite. Despite the high accuracy on tumbling satellites, the feature analysis approach was unable to differentiate clearly between stable satellites. In contrast, the differential approach was able to achieve comparable accuracy with the raw data approach but the significantly increased processing time and volume of data does not provide any significant advantage against the raw data approach. High accuracy for attitude mode identification is common across all approaches, differentiating between tumbling and attitude stable satellites.

The real data testing showed that it is possible to train a neural network with simulated data and then apply this network to real data with accuracy of $70 \pm 10\%$ only if the simulated data contains the same levels of gaps and noise as is present in the real data. The fluctuation in the real testing accuracy is likely caused by the limited real data (only 10 examples per satellite, all in the same month) compared to the thousands of training examples that can be simulated, each with a randomised configuration. A dedicated campaign to supplement the existing photometry would be a logical next step in continuing this work.

6 ACKNOWLEDGEMENTS

This work is still ongoing funded by the Defence Science and Technology Laboratory (Dstl) under the contract DSTLX-1000155513.

7 REFERENCES

- Emmons, R.H., Rogers Jr, C.L., & Preski, R.J. (1967). Photometric Observations of Artificial Satellites for Determining Optical and Physical Characteristics. *The Astronomical Journal* **72**(8), 939-944.
- Linder, E., Silha, J., Schildknecht, T., & Hager, M. (2015). Extraction of Spin Periods of Space Debris from Optical Light Curves. In *66th International Astronautical Congress*. IAC-15-A6.1.2. Curran Associates, Inc.
- Bradley, B., & Axelrad, P. (2014). Lightcurve inversion for shape estimation of geo objects from space-based sensors. In *International Space Symposium for Flight Dynamics*. Univ. of Colorado
- Furfaro, R., Linares, R., & Reddy, V. (2019). Shape identification of space objects via light curve inversion using deep learning models. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Schill, S. R., Jensen, J. R., Raber, G. T., & Porter, D. E. (2004). Temporal modeling of bidirectional reflection distribution function (BRDF) in coastal vegetation. *GIScience & Remote Sensing*, **41**(2), 116-135.

6. Kurihara, T. and Takaki, Y. (2012). Shading of a computer-generated hologram by zone plate modulation. *Optics express*, **20**(4), 3529-3540.
7. Akenine-Möller, T., Haines, E. and Hoffman, N., (2019). *Real-time rendering*. Crc Press.
8. Seidelmann, P.K. ed., 2006. *Explanatory supplement to the astronomical almanac*. University Science Books.
9. Cognion, R.L. (2013). Observations and modeling of GEO satellites at large phase angles. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*.
10. Mahabal, A., Sheth, K., Gieseke, F., Pai, A., Djorgovski, S.G., Drake, A.J. and Graham, M.J. (2017). Deep-learnt classification of light curves. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE. (pp. 1-8)