AFRL-AFOSR-UK-TR-2024-0023

Understanding the optical lightcurves of LEO spacecraft: the application of machine learning techniques

Pollacco, Don
UNIVERSITY OF WARWICK
UNIVERSITY HOUSE - WEST MIDLANDS
COVENTRY, , CV4 8UW
GBR

03/08/2024
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
Air Force Office of Scientific Research
European Office of Aerospace Research and Development
Unit 4515 Box 14, APO AE 09421

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	**2. REPORT TYPE**	**3. DATES COVERED** START DATE	END DATE
20240308	Final	20180930	20230329

4. TITLE AND SUBTITLE
Understanding the optical lightcurves of LEO spacecraft: the application of machine learning techniques

5a. CONTRACT NUMBER	**5b. GRANT NUMBER**	**5c. PROGRAM ELEMENT NUMBER**
	FA9550-18-1-7017	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Don Pollacco

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	**8. PERFORMING ORGANIZATION REPORT NUMBER**
UNIVERSITY OF WARWICK UNIVERSITY HOUSE - WEST MIDLANDS COVENTRY CV4 8UW GBR	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**	**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**
EOARD UNIT 4515 APO AE 09421-4515	AFRL/AFOSR IOE	AFRL-AFOSR-UK-TR-2024-0023

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This report is officially from a 6 month period and is a duplicate of the previous years report. During this period the paper shown here was undergoing assessment by the Advances in Space Research Journal. This was a traumatic experience as the journal spent many months unable to get anyone to assess the paper. After 9 months (summer 2023) we received two referees reports but we were not happy with them as it was clear these referees were not expert in machine learning techniques and we felt the reports were not useful. Given that we consider this work somewhat ground breaking we really wanted data scientist experts to critically examine the manuscript. Consequently, in September 2023 we removed the manuscript from Advances in Space Science and resubmitted it to the RAS Techniques and Instruments (RASTI) Journal. While this is a normal space research journal it benefits from having access to real experts in these techniques. In addition the editors from RASTI have expressed interest in increasing the SDA content of their journal. Anyway in Dec 2023 we received 2 referees reports from RASTI which were clearly from experts. They are extremely supportive of the paper as it stands and have suggested an additional small piece of work. They requested an examination of MMT9 database for any debris lightcurves and to see how they score in our framework. While we have agreed to this I’m in two minds as to its usefulness. This is because the MMT9 system is composed of a number of small aperture telescopes and only the largest debris (ie satellite sized) will be detectable. It stands to reason that these will follow the same rules as the other platforms. The exception to this is satellites that have become debris due to damage. In essence we have already been looking at this as we have been examining within the current framework evolutionary changes in platform design. We had always intended that this would form the basis of a future paper once the original paper is published, so we need to find a way to placate the referee without giving too much from the future paper away!

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**
U	U	U	SAR	12

19a. NAME OF RESPONSIBLE PERSON	**19b. PHONE NUMBER** *(Include area code)*
GRANT THOMAS	314 235 6046

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

FA9550-18-1-7017

Understanding the optical lightcurves of LEO spacecraft: the application of machine learning techniques

Don Pollacco, Billy Shrive & Paul Chote

Period: 30 SEP 2018 – 29 Mar 2023

Table of Contents

Summary

Additional summary added ~January 2024
This report is officially from a 6 month period and is a duplicate of the previous years report. During this period the paper shown here was undergoing assessment by the Advances in Space Research Journal. This was a traumatic experience as the journal spent many months unable to get anyone to assess the paper. After 9 months (summer 2023) we received two referees reports but we were not happy with them as it was clear these referees were not expert in machine learning techniques and we felt the reports were not useful. Given that we consider this work somewhat ground breaking we really wanted data scientist experts to critically examine the manuscript. Consequently, in September 2023 we removed the manuscript from Advances in Space Science and resubmitted it to the RAS Techniques and Instruments (RASTI) Journal. While this is a normal space research journal it benefits from having access to real experts in these techniques. In addition the editors from RASTI have expressed interest in increasing the SDA content of their journal. Anyway in Dec 2023 we received 2 referees reports from RASTI which were clearly from experts. They are extremely supportive of the paper as it stands and have suggested an additional small piece of work. They requested an examination of MMT9 database for any debris lightcurves and to see how they score in our framework. While we have agreed to this I’m in two minds as to its usefulness. This is because the MMT9 system is composed of a number of small aperture telescopes and only the largest debris (ie satellite sized) will be detectable. It stands to reason that these will follow the same rules as the other platforms. The exception to this is satellites that have become debris due to damage. In essence we have already been looking at this as we have been examining within the current framework evolutionary changes in platform design. We had always intended that this would form the basis of a future paper once the original paper is published, so we need to find a way to placate the referee without giving too much from the future paper away!

Regarding our contract with USAF/Eaords (FA9550-18-1-7017) the simplest solution as we move forward beyond the original contract then additional papers that come from this work will be sent to the Eaord office for examination. We have several more almost ready to be submitted once the first is finalised.

This is a version of the original report.
During the Sept 2021-Sept 2022 period we have been applying relatively simple decision tree analysis to our LEO database (supplemented with mmt9 lightcurves). We found that this was a promising approach and found good classification (88-98%) when tested against a trained dataset. Classification for

payload/rocket body was not quite as accurate (76%) but still useful. Classification clustering was searched for with a self-organising map analysis but the results from this were of low significance. Preliminary analysis again using decision trees of some LEO platforms showed surprisingly good classification potential. This work is ongoing.

Introduction

LEO orbits are now being extensively used now that we have entered the era of the so called “mega-constellations”. Low orbits filled with swarms of spacecraft can be used to gain some of the advantages of GEO orbits but with decreased latency. Hence, projects using LEO constellations are aiming enable cost effective internet services over the complete globe and will be heavily used by the financial industry. Amongst the thousands of new satellites there will be opportunities for new actors to hide or at least disguise new spacecraft and their instrumentation. The FA9550 experiment is based around obtaining high quality and high time resolution light curves of LEO objects and then using sophisticated image analysis and signal processing techniques to try and probe the nature of the body and to understand what confidence can be placed on these inferences.

Our work in FA9550 has its rootes in our developments of astronomical work, specifically in the assessment of shallow eclipses seen in Kepler lightcurves and determining the likelihood of these being due to an exoplanet or, more likely, due to some other astronomical mimic/phenomena. In this work we learnt the importance of using the highest quality data to help minimize systematic instrumentation effects (which could resemble astronomical signals). It is better to eliminate these effects from our data rather than try and correct for them later.

In the case of LEO targets we have endeavoured to understand the systematics produced as an object moves rapidly through the sky and the impact this has on measurements of position, time and brightness. Of course the potential for systematics are much greater for moving objects so we spent sometime trying to understand and identify these effects early on in this project.

For this experiment we used the SuperWASP lens array on La Palma (Pollacco et al 2006). This facility is an ultra-wide field imaging camera composed of 8 different optical tube assemblies (OTA). The OTA configuration can be manually changed to allow different amounts of field overlap. SuperWASP was designed from the start to facilitate extremely accurate photometric brightness estimates (eg the lens elements in the OTAs are heated to a constant temperature all year round to eliminate focus changes). While there are still instrumental based systematics, in general, we do understand the instrument in detail. For example, during times of rapid environmental temperature changes (eg a few hours after sunset) focus changes in the OTA can induce image profile variations, which intern can produce measured brightness variations. We solved this by heating the lenses in the OTAs to 34C constantly throughout the year so while there are still profile variations (now caused by gravity acting on the tubes) the overall impact is small.

For FA9550 we used data obtained from SuperWASP but also found data from the Mini-MegaTORTORA (MMT-9; Karpov et al 2015) project useful for further tests.

Methods, Assumptions and Procedures

We had always arranged that the work for this proposal would be undertaken by a PhD student but as the FA9550 was awarded late in the academic year (30th September 2018), the student we had identified to start on 1st October 2018 was unable to commit to the project and moved into industry. Given this late loss we decided to recruit a student with a view to start sometime in 2019. After an open recruitment process we appointed Arthur Greenwood (awarded a first class Physics degree from Southampton University) which although he had no formal machine learning experience, he was enthusiastic and had a clear track record of data analysis. Arthur commenced his work in August 2019. Unfortunately, Arthur was unduly effected when the COVID-19 pandemic arrived and he abruptly withdrew to return home at the end of March 2020. During this period we continued to take LEO data from SuperWASP and optimized the data reduction. By summer 2020 we had collected some 2500 lightcurves – this would form the dataset for our subsequent analysis.

Over 2020, we identified an additional student, Billy Shrive, who was completing his Mathematics and Physics degree at Warwick. His final year project involved significant data modelling with sophisticated algorithms. Billy has demonstrated ability in mathematics and computational physics. Billy started on this project in August 2020 and at the time of this report has had 13 months in position.

All the above was pointed out in previous end of year FA9550 reports.

Within 1 month of starting Billy Shrive was experimenting with a number of well known machine learning techniques (mostly with software he had written himself):

1. K-nearest neighbours
2. Support vector machines
3. Decision trees
4. Neural networks

We identified that decision trees could be particularly valuable for our LEO analysis. These are used to predict the value of a target variable using multiple input variables (such as predicting one property of space debris based on others). Decision trees separate input variables into subsets to predict the outcome of the target variable. The resulting tree creates a set of rules to move through, such that any new datapoint can be used to determine the corresponding value of the target variable. The target variable can take discrete values (here the tree is called a classification tree) or continuous values (regression trees).

During this year Billy gained experience with the SuperWASP and Mini-MegaTORTORA data as well as applying machine learning techniques to the lightcurves.

Results and Discussion

During the report period we finalied the data reduction (ie re-reduced all the SuperWASP data again) to produce our final data products. In total we obtained 2500 lightcurves from around 950 objects, selection was based mainly on visibility and expected brightness (the instrument is sensitive to targets in 6th-10th Magnitude range and biased towards objects in higher LEO orbits which are visible for longer during the night (Chote et al 2019).

While lightcurves were given in the previous yearly report we include them again here as this work encompassed both periods.
Figures 1-3 show the lightcurves of a range of objects to illustrate the quality that can be obtained using the techniques we have developed. One of the features of this dataset are the gaps in the lightcurves caused by the telescope moving to a new mount position.

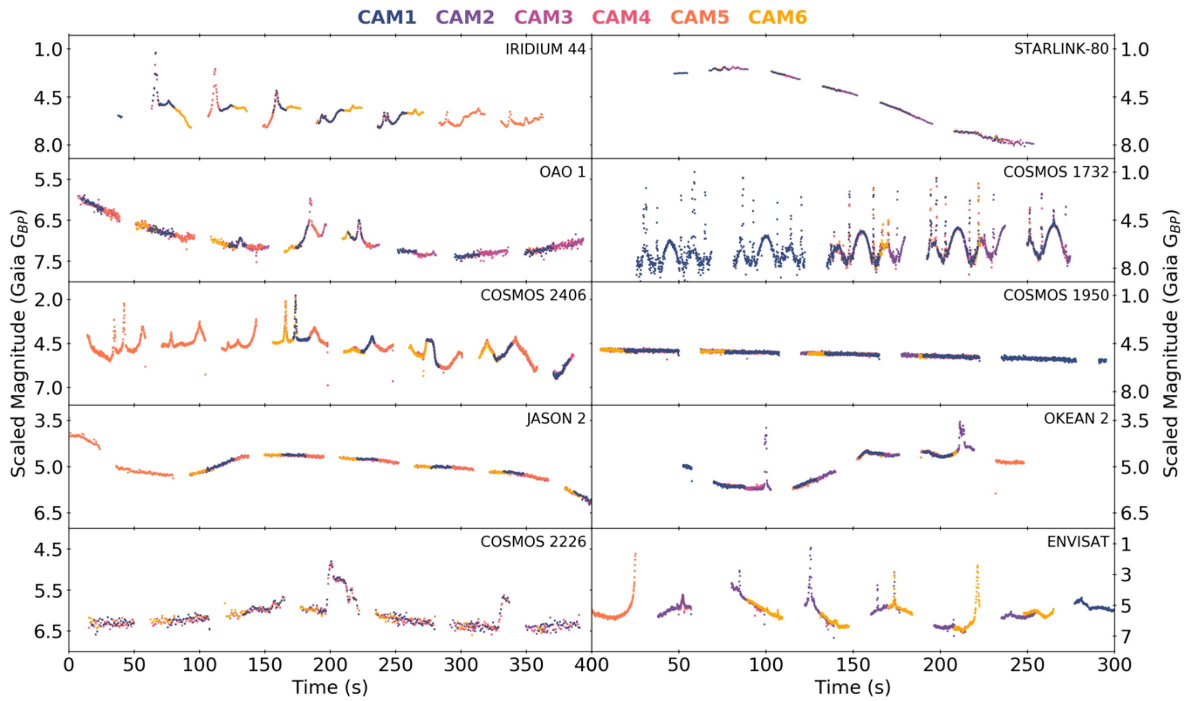

Figure 1: Example Satellite lightcurves from the LEO survey

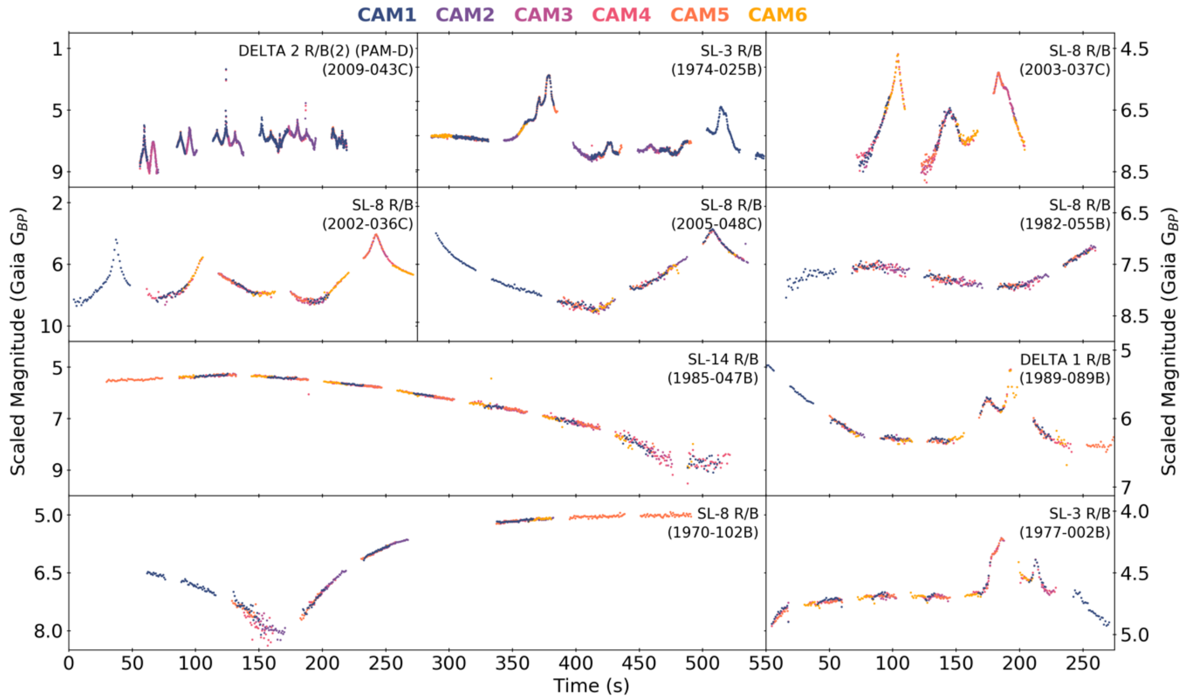

Figure 2: example rocket body lightcurves from the LEO survey

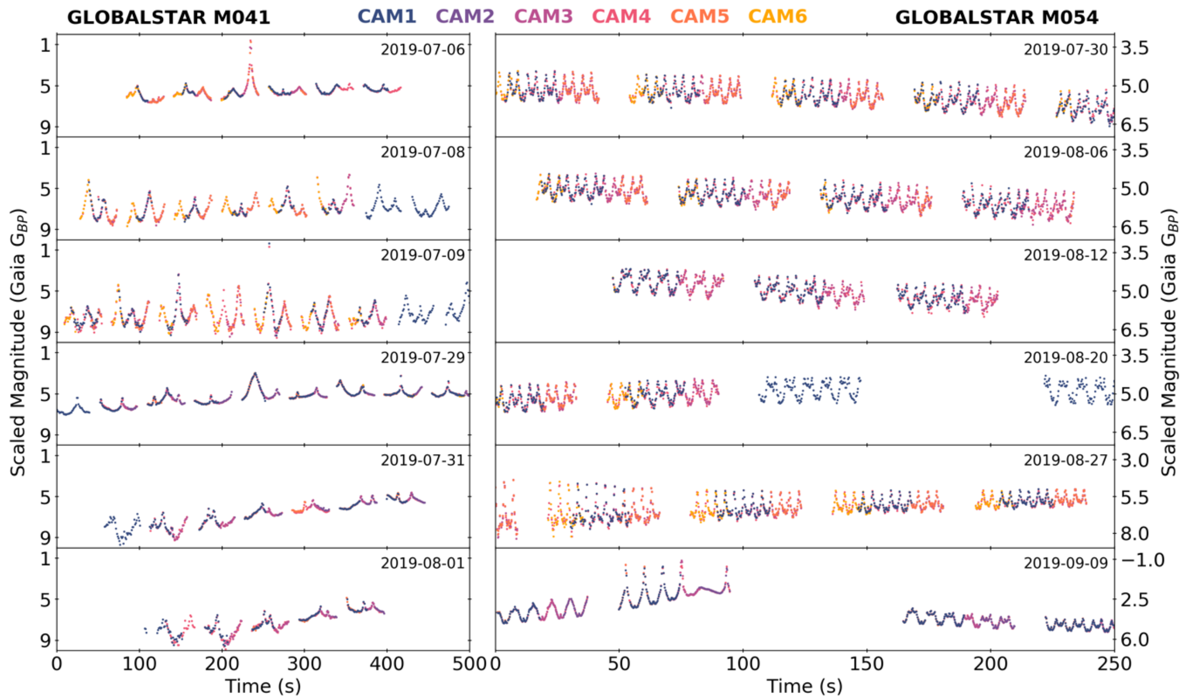

Figure 3: Example repeated light curves showing the time-evolution of two defunct satellites from the Globalstar constellation that were boosted into graveyard orbits at the end of their service life.

Decision trees and Self Organising Map analysis – initial analysis

It became clear early on that for classification tasks with a small number of classes, decision trees will outperform a Deep Neural Network, especially with respect to speed (several magnitudes faster to reach similar accuracy). Therefore, we decided to continue the analysis with decision trees.

There are two main types of algorithms used, using either Gini impurity or information gain to determine the subset boundaries. Gini impurity is a metric of the frequency of incorrect labelling, if the labelling was random (with an accurate distribution of labels). The Gini impurity for a dataset with J classes, with p_i the fraction of items in class i, is given by Eqn. 1.
Information gain is a term in information theory, meaning the amount of information gained about random variable X from observing a different random variable A. Information gain is related to entropy, being equal to the entropy of the parent variable minus the sum of the entropy of the children variables (with the children variables represented as a) (Eqn. 2).
Using more than one decision tree is a technique used in several methods. Boosted trees is one such method, building an ensemble of decision trees with each new tree giving consideration to the previous instances. An ensemble of decision trees can also be used in bootstrap aggregated decision trees, where each tree is trained by resampling training data and using all the trees at once to form a consensus.

From our LEO database we used 2240 lightcurves, 60 of these omitted (bad data or missing type) to leave 2180. The lightcurves stored in JSON format, for which we extracted times and magnitudes along with the type of satellite. Using the 'feets' python module, 55 statistics were calculated for each lightcurve, at 3 binning levels, resulting in 165 statistics. These covered a range of parameters and we manually labelled the lightcurves as smooth varying, periodic and type (rocket body/payload/other). The dataset was split into 380 lightcurve for training and 1800 for testing,

Based on the labels marked previously (smoothness, periodicity and satellite type), a boosted decision tree classifier was trained to predict these labels. With regard to satellite type, >95% of the satellite were classed as either 'payload' or 'rocket body', so this was the binary classification chosen for the decision trees.

After one minute of training, the resulting models plateaued, getting the follow percentage of predictions correct:

- Smoothness: 98% accuracy
- Periodicity: 88% accuracy
- Type: 76% accuracy

As is usually done we computed a receiver operating characteristic (ROC) curve for each of these models, comparing the trade-off between false positives and true positives (Figure 4). The ROC curve above the diagonal line indicates good classification (perfect classification would encompass the point (0,1)).

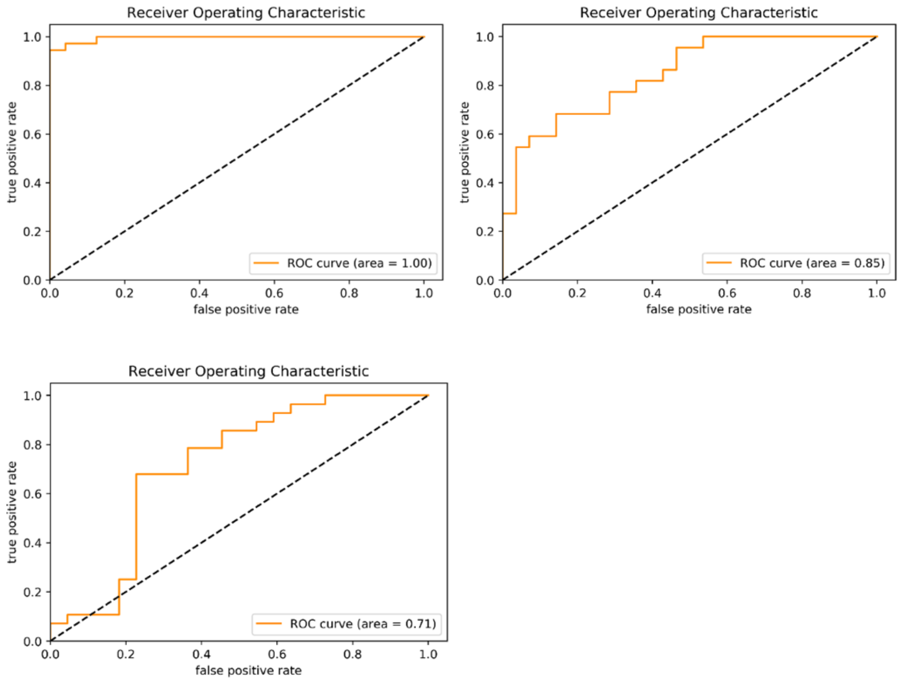

Figure 4: Decision tree ROC curves: top left - smoothness (very good classification), top right – periodicity (good classification), bottom left – smoothness (reasonable classification)

The contribution of each statistic to the decision trees was used as a metric for the usefulness of that statistic. The top N useful statistics were then used to generate a self-organising map (SOM) for each classification task. N=100 was found to generate the best clustering for each task (see Figure 5 – left panel).

While the clustering in this SOM was clearly low we noted that the dominant statistics seemed to often involve frequency and hence related to periodicity. However, the SOM for satellite type showed almost no evidence for clustering, but we suspected that increasing the number and variety of statistics used was worth testing to see if clustering could be revealed. Hence, we generated 785 different statistics for each lightcurve. For various N, no improvements could be made to periodicity, but for satellite type there were improvements to the SOM, with N=500 producing the best results (Figure 5 – right panel).

After some deliberation we concluded that given the weak clustering in all the SOMs and the difficulty in interpreting what the statistics actually meant, we decided that the decision tree approach contained more useable information and we would progress that approach.

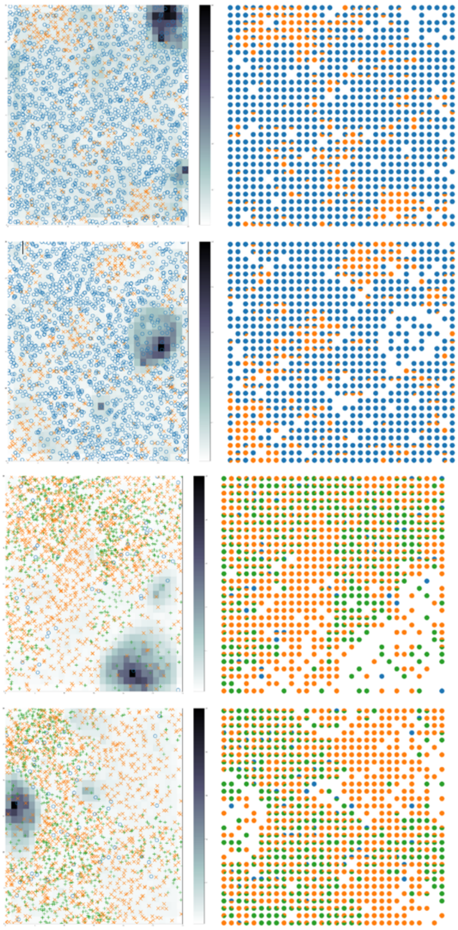

Figure 5: Left panel: two example SOMs of periodicity (orange) vs non-periodic (blue) determined from the 100 statistics. Right panel: two example SOMs of satellite type; debris (blue), payload (orange), rocket body (green) determined from 500 statistics.

Decision trees – further analysis

Before progressing our analysis we would do some further tests on our LEO data. Just to recap: the SuperWASP data is obtained by moving the mount along the satellite TLE track and exploits the wide field of the instrument. The exposure length is calculated from a position along the track when the satellite is within the field of view and ends before the satellite leaves the imaging field. Consequently, the start and end times (and positions) of the satellite in the field are extremely well defined and the light curve generated by interpreting between these points using the TLE to generate the object the expected velocity at any particular time. A feature of this is that the lightcurves are *almost* randomly interrupted by gaps as the telescope mount moves to its new position. Initially we noted than many seemed to have similar gaps but on close examination they were all slightly different – as you would expect for objects passing the observatory on different trajectories.

Anyway, given that these lightcurves are “different” to those usually seen we thought it best to get a different set of “normal” lightcurves (ie no gaps) and do a similar decision ree analysis.

We retrieved several hundred light curves of various LEO targets from mmt9 for the analysis. We used these to produce a second series of lightcurves, identical to the originals but with data gaps inserted (reproducing the properties of gaps seen in the SuperWASP data). We then repeated the same analysis for both series. Reassuringly the decision tree results were almost identical between the gapped and continuous data but the gapped data showed slightly less confidence (1-2% less). While this result was reassuring it is to be expected. We also noticed the vast quantity of data available within the mmt9 database and are continuing to use it.

We then decided to extend the decision tree analysis to see what it could tell us about satellite platforms. We extracted 100 lightcurves of each of the Starlink, Iridium-plus and Globalstar platforms, using 80 of each for training and 20 for testing. We were surprised to find that the decision tree analysis was able to identify lightcurves with good certainty from each platform.

This must indicate they the spacecraft themselves are physically distinction in some way. Currently, we are examining images of the spacecraft and our speculative thoughts are that the Starlink spacecraft have a single sail and are clearly quite different to the others. This is less true for the Iridium-plus and Globalstar satellites which do resemble each other. However, closer examination shows they have quite distinctive as their different sensor suites are positioned at one end of the satellite – the end that would be facing the observer on each transit. We have also compared the untrained data for each platform against all other (trained) platform data with the results given in Table 1.

Table 1: Recovery rate from decision tree analysis for three major satellite platforms.

Platform (trained)	Unclassified Test Data		
	Starlink	Iridium-plus	Globalstar
Starlink	18	2	3
Iridium-plus	0	17	1
Globalstar	2	1	16

This, at first glance is rather encouraging: not only can we identify lightcurves of each platform within their platform with high confidence, we can also tell that they are highly unlikely to belong to another platform. These results are all very surprising but we have high confidence in their validity (at least for these platforms).

On-going and Future Work

As the results of the decision tree analysis are really quite surprising, we want to extend this parameter space ie complete the same analysis for additional platforms. As we extend this analysis we will have to examine