HyperClass - Executive Summary

Massimiliano Vasile[1], Simão da Graça Marto[1], Lewis Walker[1], Andrew Campbell[2], Vasili Savitski[3], Paul Murray[2], Stephen Marshall[2]

July 18, 2023

[1]Department of Mechanical and Aerospace Engineering, University of Strathclyde
[2]Department of Electronic and Electrical Engineering, University of Strathclyde
[3]Fraunhofer UK Centre for Applied Photonics

Project Scope and Objectives

Project HyperClass aimed at accelerating the technology to process spectra received from single pixel images of space objects and: i) extract the material composition of the object, ii) classify the object according to a set of predefined features, iii) estimate the attitude motion from the spectral lightcurve.

The specific objectives of project HyperClass were:

1. To extend the current high fidelity model and simulation to a wider range of materials and material properties, material distributions, geometries, attitude motions, optics.
2. To demonstrate, in simulation, the applicability of hyperspectral light-curve analysis to a wide range of objects with different surface composition, size, shape, attitude motion, orbit regimes, illumination conditions, in space and on ground sensors.
3. To demonstrate the use of deep learning for the classification of space objects from hyperspectral images and hyperspectral light curve analysis.
4. To demonstrate the use of deep learning for attitude motion reconstruction from hyperspectral light curve analysis.
5. To associate hyperspectral light curve analysis and attitude motion to identify patterns in space object behaviour
6. To design a prototype sensor that can be used in conjunction with a standard telescope on ground or in space on board a satellites
7. To test the concept in a lab environment with a mock-up of a small satellites. Possibly to be revised to: testing the concept with actual optical observations from ground.
8. To demonstrate the use of multispectral/hyperspectral imaging to identify the surface composition of resident objects in Low Earth Orbit, either using on ground or in orbit observations, and to improve their classification.
9. To demonstrate that the use of the time variation of the intensity at different wavelengths can be used to reconstruct the attitude motion better than by simply using light curve measurements.

During the development of HyperClass we achieved all the objectives listed above. However, we found that some results could be achieved with machine learning but not necessarily deep learning. The use of deep learning and the possible advantages of deep learning are a possible extension for future work. Furthermore, we included a part on material ageing that was not initially foreseen but that has produced quite interesting results. We now have a high fidelity simulator that can incorporate a wide range of objects, materials, material ageing, attitude motions and spectra including specular reflections.

Main Results

HyperClass achieved two main results: i) the robust classification of space objects from single pixel observations via spectral unmixing and material library matching, ii) the estimation of the attitude from the spectral light curve analysis. The former was tested both in a lab environment and with images of real space objects acquired with a telescope mounted on the roof of one of the buildings at Strathclyde. The latter was tested in a lab environment with real image sequences and in simulation using our high fidelity simulator.

Probabilistic Object Classification

One of the key outcomes of HyperClass is to classify satellites into classes that describe their general composition or intended purpose. We used the probability of each material being present on the object to classify the object (see Figure 2.1). The probability of detection was achieved with both a classification based on a least square approach using a material library and a machine learning classifier based on the same library. In both cases the spectrum is acquired and unmixed using the spectral library. After unmixing the residual, or difference between the spectrum generated with the spectral library and the actual spectrum is retained for further analysis.

Satellite	Classification
Starlink	Comms/GNSS
Iridium-NEXT	Comms/GNSS
HS376	Comms/GNSS
Galileo	Comms/GNSS
GPS Block IIF (GPSIIF)	Comms/GNSS
GLONASS-K	Comms/GNSS
Upper Stage	Rocket Body
DubaiSat-2	EO
CALIPSO	EO
CubeSat (panels)	CubeSat
CubeSat (no panels)	CubeSat
Box-Wing	None
Dragon2	Capsule

Table 2.1: Full list of satellite models used and corresponding classes

While the abundance of material present can be detected and is theoretically also a useful predictor, this is intentionally not used in the current model. This is because the abundance of material that is measured correlates with the randomised movement of the simulated objects, i.e. measuring half the expected abundance may simply mean the component made from that material is partially obscured during observation. Therefore, in this initial classification model the abundance is not used and instead it is explored how accurate a classification could be made using only the previously derived probability of each material being present. For this trial satellites are split into the classes Comms/GNSS, Rocket Body, Cubesats, Earth Observation (EO) and Capsule as defined in Table 2.1. Note that the Boxwing satellite is a generic shape so has no obvious class.

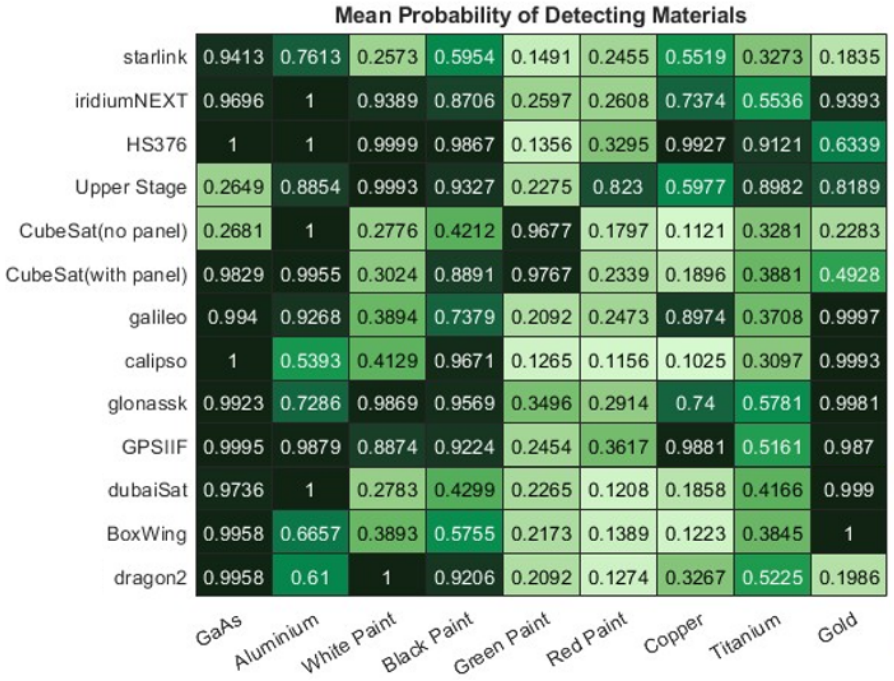

Figure 2.1: Mean probability across 2000 samples of each satellite of a material being present on that satellite according to the ANN model

As mentioned previously training a model on examples of these objects achieves high classification accuracy but also leads to over-fitting the model to the satellites in the training set. To avoid this the model is instead trained on synthetic combinations of material probabilities. This has the advantage of not using satellite data during training and allows the full range of combinations of material probabilities to be used, not just those combinations that occur in the training datasets.

Synthetic data is generated by first defining a probability range for each material in each class, that specifies the highest and lowest probability that may occur in measurements. To do this it is necessary to make assumptions about the materials in each class. This is achieved by studying

Class	GaAs	Aluminium	White	Black	Green	Red	Copper	Titanium	Gold
Comms/GNSS	90-100%	50-100%	0-100%	0-100%	0-30%	0-100%	30-100%	0-100%	0-100%
EO	90-100%	50-100%	0-30%	50-100%	0-30%	0-30%	0-30%	0-30%	50-100%
Capsule	90-100%	50-100%	50-100%	50-100%	0-30%	0-30%	0-30%	0-30%	0-30%
Rocket Body	0-30%	50-100%	50-100%	50-100%	0-30%	50-100%	0-30%	50-100%	50-100%
CubeSat	0-100%	50-100%	0-30%	0-50%	50-100%	0-30%	30-100%	0-30%	0-30%

Table 2.2: List of classes and range of material probabilities used in synthetic training dataset

the materials that are expected to be present on each satellite, as shown in Figure 2.2.

		GaAs	Al	White	Black	Green	Red	Cu	Ti	Gold
Comms / GNSS	iridiumNEXT	■	■	■				■		■
	Starlink	■	■					■		
	HS376	■	■	■	■		■	■	■	■
	GPSIIF	■	■	■	■			■		■
	Glonassk	■	■	■	■			■		■
	Galileo	■	■		■			■		■
EO	DubaiSat	■	■		■					■
	Calipso	■	■		■					■
Capsule	Dragon2	■	■	■	■					
Rocket Body	upperStage		■	■	■		■		■	■
CubeSat	CubeSat (w/ Panels)	■	■			■		■		
	CubeSat (w/o panels)		■			■		■		
None	Box-wing	■	■		■					■

Figure 2.2: Materials that are expected to be present on each object in the simulated dataset

For example, if it is assumed titanium is present in all rocket bodies then a suitable range for synthetic data for rocket bodies might include titanium probabilities in the range 50-100%. Meanwhile, for Comms/GNSS satellites Titanium is present in some but not all objects, so the synthetic data for the Comms/GNSS class may include titanium probabilities in the range 0-100%. When generating this data it is helpful to ensure all extrema of these ranges are included in the synthetic data. The training data used in this study is of the ranges shown in Table 2.2 and for each class every combination of material probabilities is generated with the limit that only the maximum, minimum and median value of the defined range are used. This resulted in 358,000 synthetic combinations being used in the training. Optimising this synthetic training data based on effect on the subsequent object classification model would risk the over-fitting scenario this approach intended to avoid. However, some observations about the original material detection model were used to define optimal ranges for the table in Table 2.2. It can be observed from Figure 2.1, that false positive probabilities ¿40% are likely to be rare in the proposed model. So when generating synthetic data the probability generated for non-existent material is defined so as to not exceed 30%, while when materials are present values between 50-100% are typically generated. For copper, which is harder to predict given the low abundance, the generated training data includes probabilities of prediction as low as 40%, when copper is present, to reflect this low abundance. Similarly for solar panels the generated data includes only probabilities above 90%. These threshold help generate training data but are not threshold on the model, which will learn and apply its own classification thresholds based on the relationship between these datasets.

This is an arbitrary definition of the classes but it is useful to group objects based on their composition. The scope of this definition of the classes in this report is to show an example of how to classify from the probability identification of materials and to illustrate the possible critical aspects of the classification process. Alternative definition are indeed possible.

Comms/GNSS satellites and EO satellites are clearly similar but should be distinguishable by the copper components. Similarly the paints present should distinguish Capsule and CubeSats from other classes, while the lack of GaAs should help identify rocket bodies. While the generic BoxWing satellite has no obvious classification it is most similar to those in the EO class so it would be expected that this would be the most probable classification.

A k-nearest neighbours (KNN) model is trained using the synthetic combinations of material probabilities defined in Table 2.2. The benefits of this method in this case is that classification is based on a distance in feature space between the observation and training data instances of each class. It is, therefore, well suited to classifying objects where the probabilities do not exactly match the ideal cases shown in the dataset. The prediction rate for the 2000 simulated instances of each

satellites into the each class is shown in Figure 2.3.

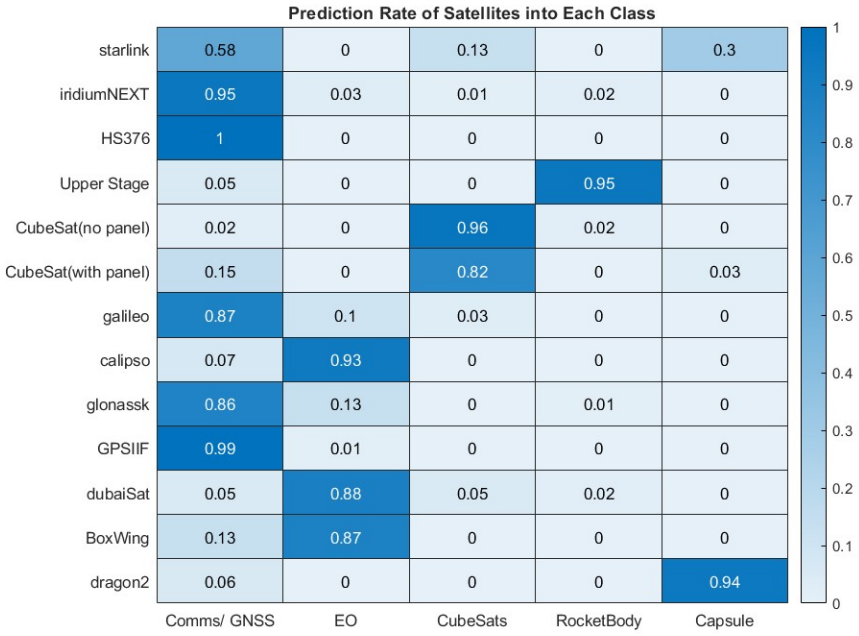

Figure 2.3: Prediction rate from the knn model when classifying objects based on probability of materials present on the object

In the majority of cases this classification model produces convincing results. Most classes are classified well and the overall classification accuracy across all satellites is 89%. Across 12/13 objects classification accuracy is 92%, as most errors are in Starlink satellites. For example, confusion between Comms/GNSS and EO is the most common error, which is to be expected given the similarity of these objects and that only small abundances of copper typically distinguish them. Conversely, it is quite rare for satellites to be misclassified as rocket bodies, which is expected given the absence of solar panels, which are typically quite significant features. Boxwing, which is a generic object with no obvious or defined class class, is predicted as EO 87% of the time, as was expected given the similarity in materials.

In the Starlink class, accuracy is 28% lower than the next least accurate class. A significant number of Starlink objects are predicted as Capsules in particular, one reason for which may be a failure to detect the copper in many cases. Copper is a challenging material to detect as it often occurs in low abundances and the decomposition model returned lower probability of copper presence then any other object known to contain it, as shown in Figure 2.1. However, when investigating the ground truth data of the observation of spectra from simulated test objects, it is found that in 37% of Starlink objects no copper was present. This occurred as the component made from copper was completely obscured in these cases and did not face the observer. Thus the errors here are explainable, and while it is probable that some errors results from visible copper that is not properly detected the most significant error is that the classification is likely based on a material that was not visible in all starlink observations, so could not possibly be detected. This highlights a potential limitation of classifying satellites based only on the material present when those materials are not guaranteed to always be visible. However, the same problem could theoretically occur in other objects where the error rates were far lower, so it is reasonable to assume the geometry and materials present in Starlink are more challenging to classify.

Similarly to earlier in the pipeline, where material presence was estimated, the probability of each classification is also interesting to study and provides more explainability to the model. The KNN algorithm classifies each combination of materials based on the distance between the observed object and all the classes identified through the training data. This means that a relative metric of the probability of the observed object belonging to each class inherently exists within the model. The average of these classification probabilities across the 2000 instances of each object type are shown in Figure 2.4.

The mean probabilities are quite similar to the actual error rate, which implies most predictions are made with high confidence. However, the usefulness of this metric alongside classification results

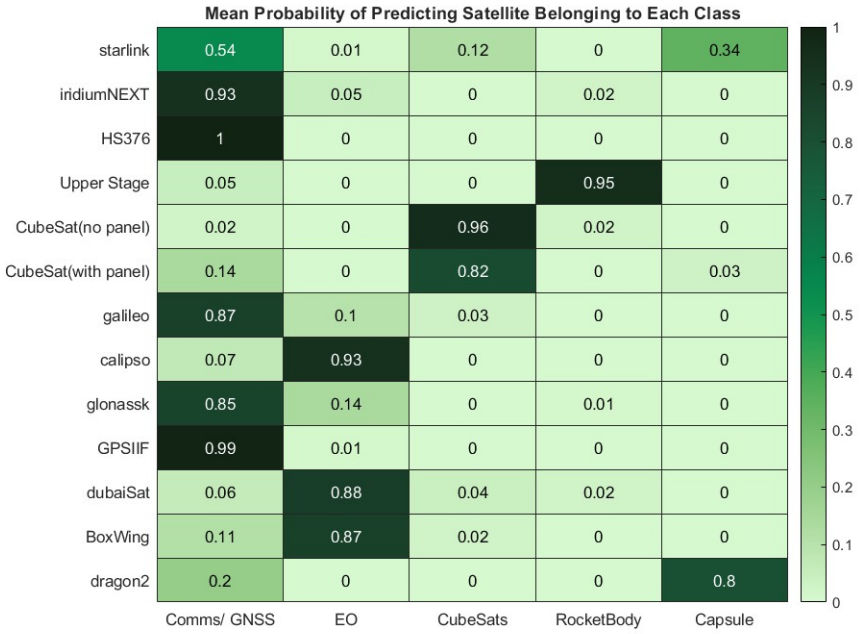

Figure 2.4: Mean probability of predicting object as belonging to each class in the KNN model

can still be observed. For example, this reveals that while 94% of Dragon2 objects were correctly identified as capsules, on average the confidence of this prediction was 80%, which is lower than the confidence of many other predictions, with a mean probability of 20% that this is a Comms/GNSS. This makes sense as the materials using in capsules are broadly similar with only a single material difference between Dragon2 and some Comms/GNSS satellites. It is encouraging that while this does not cause a 20% error rate, the similarity between these classes and potential for confusion is still captured here.

Another option for utilising probability data is to identify objects for which the correct class is highly uncertain and label these as a Unidentified Flying Object (UFO), rather than as whichever class the model deemed only marginally more probable. This is also very useful when unseen combinations of materials are present or materials not int he library are detected. Simple, manually derived checks may also be introduced here to ensure that each prediction matches key assumptions for each class. To demonstrate this theory we test the assumption that, a) solar panels (GaAs) are present in all communication satellites and never in rocket bodies and b) green paint is present only on Cubesats and not on anything else, are used as a final confidence check on the classification. Any classification that does not match these assumptions are instead labelled as UFO. The revised classification results using this method are shown in Figure 2.5. It is clear that this approach most significantly effects the Comms/GNSS class. False positive classifications in this class are almost completely eliminated, however, this occurs at the expense of failing to classify a significant number of Iridium-NEXT and Glonaask satellites, which instead are labelled as UFO. The true positive rate for this class fell from 87% to 77% using this approach while the false positive rate fell from 7% to 0.3%. Overall accuracy is lower due to the drop in true positive rate, however, the confidence in the remain classifications increases. In the proposed KNN model 91% of Comms/GNSS classifications were originally correct, but with the UFO class included this increased to 99.5%. This is a desirable behaviour because the classification system because the UFO class is, effectively, the uncertain class. Hence the classifier moves to the uncertain class all instances for which a decision has low confidence.

This provides an alternative classification approach depending on whether it is more important to identify as many objects correctly as possible or if is more important to ensure those classified are correct, at the expense of failing to classify some at all. The probabalistic approach to classification has other practical benefits not explored in this study. For example, the validation trials in this work assume a single observation of each object and no other pieces of information on orbit or motion. In practise, it is likely that objects classed with lower probabilities can be re-examined on multiple occasions to increase the information available and improve the confidence of the classification result.

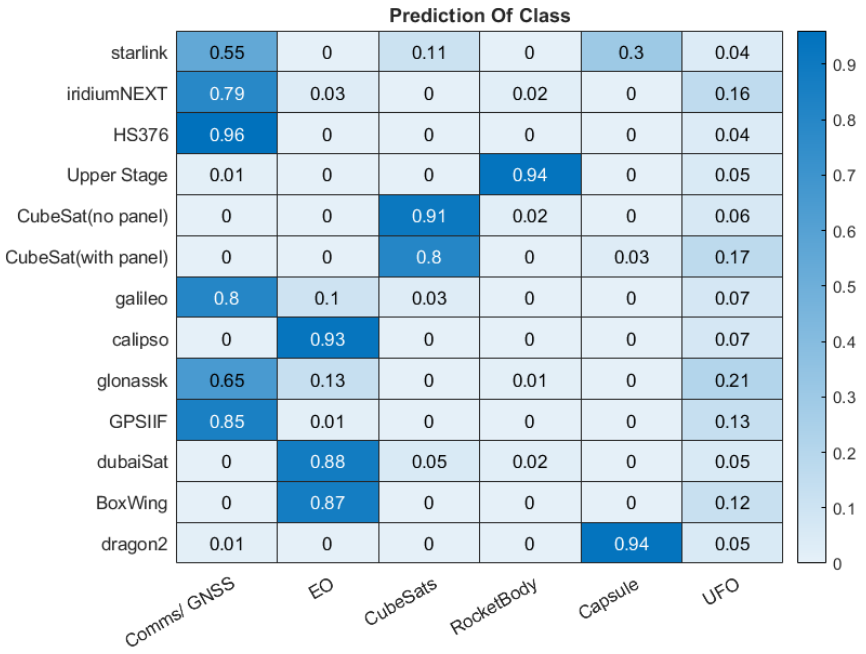

Figure 2.5: Prediction rate from the knn model when using model confidence and prior assumptions to flag uncertain predictions as UFO

Hyperspectral Attitude Estimation

The work on attitude reconstruction from spectral light curves focused on three aspects:

- Attitude angles with respect to an inertial reference frame from a single spectrum
- Attitude motion from time series of spectra - called spectral light curve inversion
- Identification of symmetry groups for spectral light curve inversion

The last point is quite important because the attitude is reconstructed from single pixel time series and with no assumption on shape or dynamics. We demonstrated that under these working assumptions even for non-symmetric objects with non-symmetric distributions of materials there are symmetry groups that would lead to the same attitude estimation from spectral light curves.

Attitude Angles from Single Observation

We considered two methods. One based on a least square minimisation to match the observed spectra with a possible combination of library spectra. The approach is very similar to the one used in TN1 to unmix the spectra but accounts also for the exact illumination conditions and view factor. The approach used an equivalent cube and aimed at determining the direction of the normal vectors to each of the faces. It was found that from a single spectrum the exact attitude could not be uniquely determined without starting from a good initial guess. The ambiguity in the identification of the exact direction of the normal vectors still allows identifying the frequency at which a particular set of materials appears in the field of view after spectral unmixing.

However, the ambiguity on the exact direction of the normals implies that the characterisation of each equivalent surface associated to a normal vector requires more than a single observation.

Attitude Motion from Spectral Time Series

Also in the case of spectral time series (or spectral light curves), it was found that the definition of the rotation vector and attitude quaternions is not unique and the right attitude and angular velocity require a good initial guess to be correctly identified.

The search for the attitude motion starts with the generation of a first guess and then a search for a minimum of a regularised loss function between observed and expected spectra. This search was performed either with an interior point method or with Differential Dynamic Programming.

The initial guess used a Phase Dispersion Minimisation algorithm plus a grid search. When considering the symmetries in the rotations and measured spectra, the method converges to one of

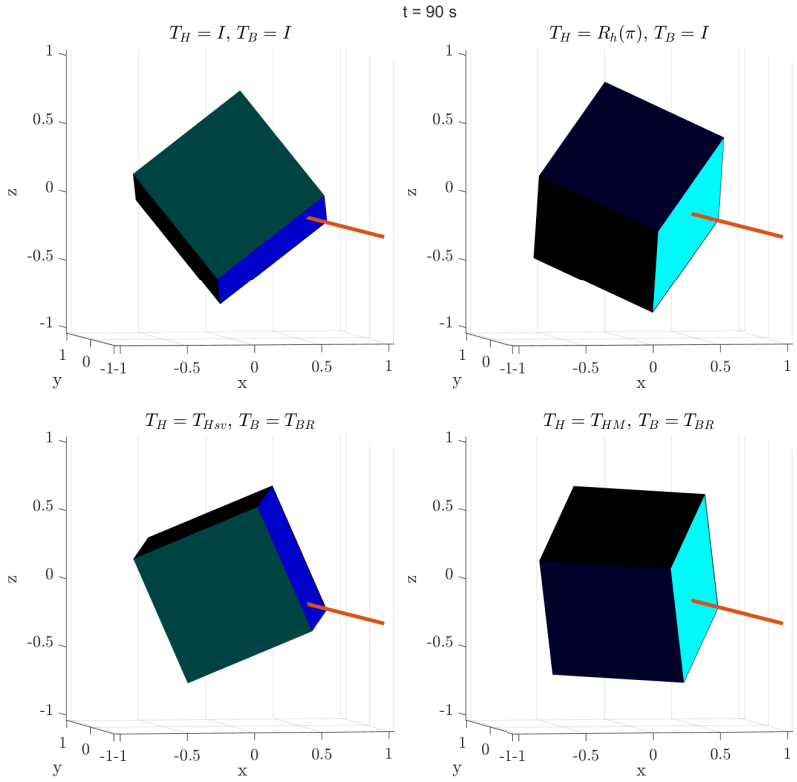

Figure 2.6: Appearance of cube object for attitudes in $\mathcal{R}$ at $t = 90s$

the expected attitudes. Figures 2.6 shows an example of converged attitude with an equivalent cube with all different faces covered in different materials. Figure 2.7 shows the correct and estimated observations.

Besides the use of the regularised least square minimisation, we tested a machine learning approach to estimate the angular velocity vector from spectral time series. The initial results are promising but require further work.

Symmetry Groups

An interesting result of the analysis of the attitude from spectral light curves is the identification of symmetries given by rigid rotation or orthogonal transformations that lead to the same spectral time history for different attitude motions. Symmetries explain also the non uniqueness of the solution for single spectrum observations.

It was found that rotations around the phase vector produce the same spectrum and spectral

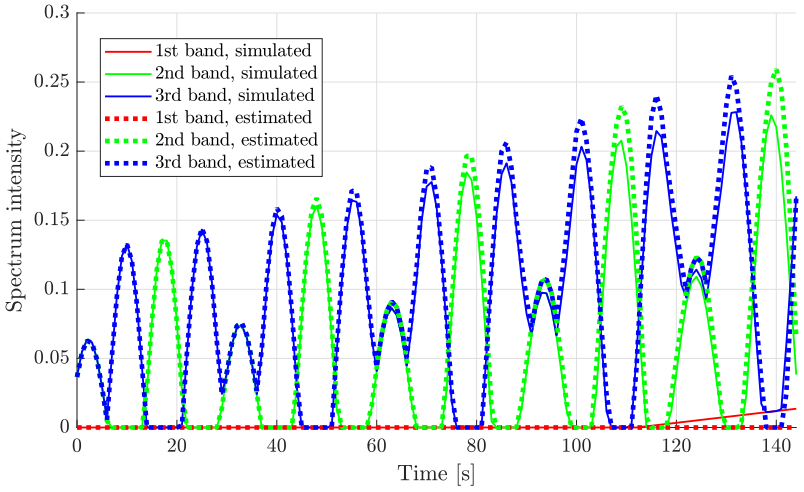

Figure 2.7: Simulated observation and the spectra resulting from the estimated attitude history

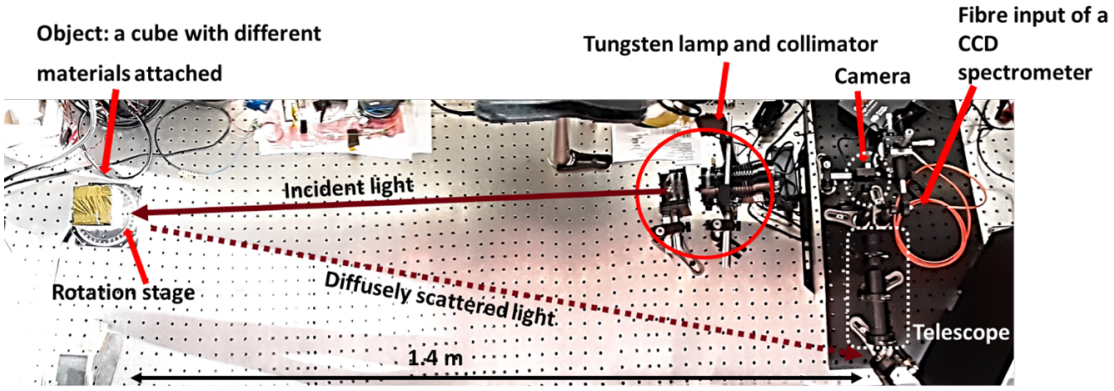

Figure 2.8: The lab prototype of the hyperspectral surveillance system assembled in Fh-CAP

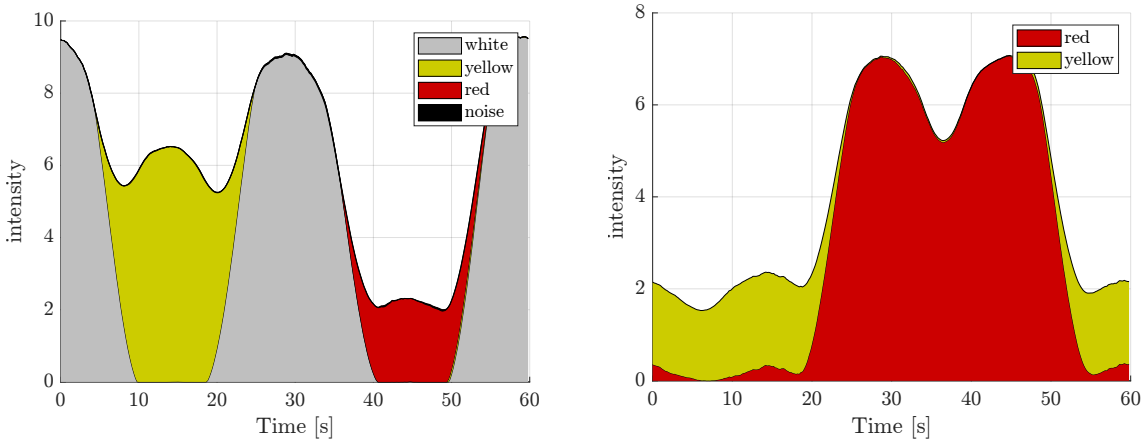

(a) Cube with faces painted in sequence white-yellow-red-white

(b) Cube with faces painted in sequence yellow-yellow-red-red

Figure 2.9: Unmixed spectra for lab data of painted cubes

time history. The continuity requirement on the attitude time history implies that the symmetry transformation has to be constant during the observation period. This limits the ambiguity on the possible attitude to a finite set of symmetric configurations plus a set of rotations around the view direction.

Experimental Testing

The concept behind HypereClass was tested in a laboratory environment to extract the correct spectra from relevant objects under relevant illumination conditions. The spectra of several materials were extracted and analysed under different illumination angles and intensities. This initial campaign to collect spectra was used to form the library used in the high fidelity simulations and in the material classification algorithms. Figure 2.8 shows the laboratory set up comprasing the illuminator, the sensor and the rotating cube.

Then the attitude estimation algorithms were tested in a laboratory by observing a rotating cube covered in known materials and under known illumination conditions. As can be seen from the comparison of the actual material distribution over time in Figure 2.10b vs the reconstructed material distribution over time in Figure 2.10b the material extraction was successfully achieved and the algorithm correctly estimated the axis of rotation.

The second set of tests were conducted by acquiring spectra of real objects in space using a telescope on the roof of one of the building at Strathclyde. With the telescope we acquired a number of objects over different nights: starlink, ISS, Cosmos and some rocket bodies. We also run a calibration test by acquiring the spectra of a solar panel of a cubesat that we positioned in Glasgow far away from the telescope. The result of the acquisition of the solar array can be seen in Figure 2.11 where the acquired spectrum is compared estimated with the spectral library.

The unmixing of the starlink and of the Cosmos satellites gave mixed results as can be seen in

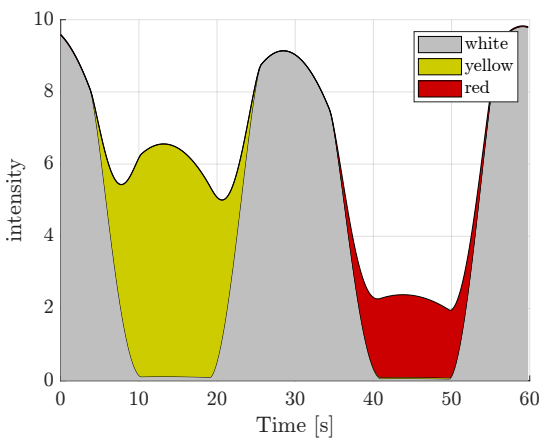

(a) Cube with faces painted in sequence white-yellow-red-white

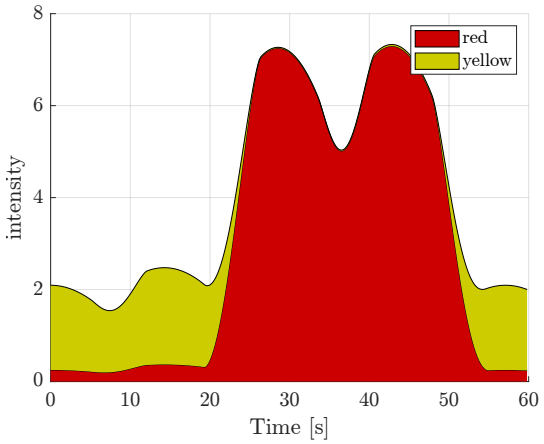

(b) Cube with faces painted in sequence yellow-yellow-red-red

Figure 2.10: Reconstructed spectra with known fixed axis rotation.

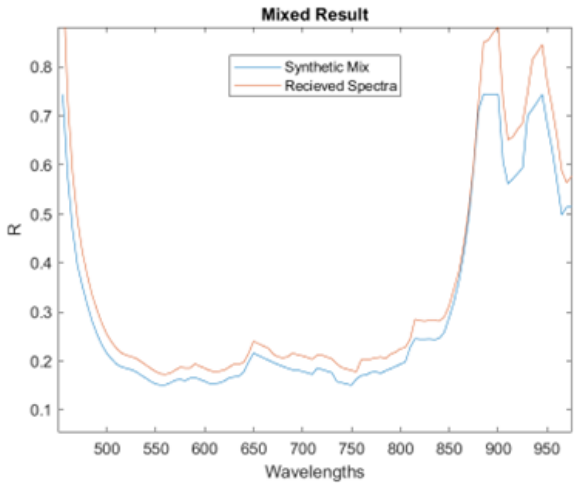

Figure 2.11: Comparison of synthetic mixed spectra (calculated from material unmixing results) and the received spectra for a known solar panel when image on earth

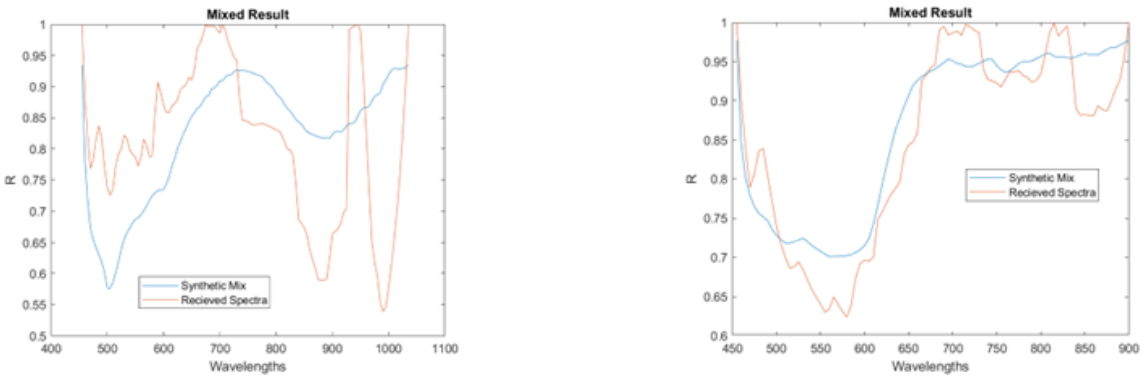

(a) (b)

Figure 2.12: Comparison of synthetic mixed spectra and received spectra: a) Starlink satellite and b) COSMOS satellite .

Figure 2.12. While the synthetic mix using the spectra in our library capture the general trend of the spectrum of Cosmos, the spectrum of starlink was quite different from even the best fit. This would suggest that our limited library with less than 15 different materials is not enough to capture the spectra received from starlink. Further work is required to validate the unmixing with known objects. It is however interesting to note that the two satellites have a substantially different spectral signature, confirming the value of using hyperspectral imaging as opposed to simple light curve or even simple colour coding.

Future Work

HyperClass achieved all the intended objectives. However, more work is required to make this technology operations. Furthermore, during the development of the project new ideas and research questions emerged that are worth pursuing in the near future. What follows is a shortlist of possible further developments.

- Extension of the material library to improve object classification, including specular reflections.
- Classification from a combination of spectral time series, spectral unmixing and estimated orbit and attitude motion
- Modelling of material ageing in space
- Testing of alternative deep learning technologies for classification - classification of components from material distribution
- Extended observation campaign from ground and comparison against known objects
- Unmixing considering the symmetry groups
- Testing of alternative machine learning technology for attitude motion estimation - extended dataset of objects
- Attitude estimation including glints and specular reflections
- Testing of deep learning technologies for attitude estimation from spectral lightcurves
- Attitude from low resolution images where the object is represented by more than one pixel - partially resolved shape of the object.
- Pose estimation from low resolution images with partially resolved object - with application to proximity operations
- Behavioural analysis of space objects
- Proximity navigation with multispectral and infrared cameras