



OPEN

Space object identification and classification from hyperspectral material analysis

Massimiliano Vasile^{1✉}, Lewis Walker¹, Andrew Campbell², Simão Marto¹, Paul Murray², Stephen Marshall² & Vasili Savitski³

This paper presents a data processing pipeline designed to extract information from the hyperspectral signature of unknown space objects. The methodology proposed in this paper determines the material composition of space objects from single pixel images. Two techniques are used for material identification and classification: one based on machine learning and the other based on a least square match with a library of known spectra. From this information, a supervised machine learning algorithm is used to classify the object into one of several categories based on the detection of materials on the object. The behaviour of the material classification methods is investigated under non-ideal circumstances, to determine the effect of weathered materials, and the behaviour when the training library is missing a material that is present in the object being observed. Finally the paper will present some preliminary results on the identification and classification of space objects.

The composition of space objects can be characterised based on the spectra of light reflecting off these objects, as each material present will commonly give unique light spectra. In the last 20 years, various techniques based on spectroscopy, spectral analysis and colour photometry have been proposed¹⁻⁷. However, the work presented in these references is mainly observational. Spectra were observed for known sets of objects and these spectra were then associated to known material types.

Colour photometry³ has been proposed to differentiate between different objects but was not used to develop a classification system. Meanwhile techniques using colour indexes, alongside characteristics of the spectra, have been used to develop a classification system, see for example^{4,6}. Both colour indexing and photometry produce relatively sparse spectral data in comparison to techniques such as hyperspectral imaging and spectrometry. In recent years, both of these spectrally rich techniques have proven useful in the classification and characterisation of asteroids⁸. Multispectral imaging has also proven useful for close proximity navigation⁹ but mainly using the near infrared part of the spectrum and for Earth Observation and remote sensing application, such as the identification of targets on the ground in¹⁰.

Classification of space objects has also been achieved using temporal rather than spectral information using light curve analysis¹¹⁻¹⁶. Most recently, Machine Learning (ML) was introduced to classify objects directly from light curves^{17,18}. As no spectral information is present these techniques rely on temporal information to reconstruct the attitude motion or the shape of space objects. As such while classification of known objects is possible, no information on the composition of unknown objects is obtained.

This paper will present methods for identifying space objects from time-varying hyperspectral sensor data using a pipeline that begins by identifying the materials present and then builds an understanding of the object before reaching a final classification. Specifically, it is proposed to first use spectral decomposition techniques to identify the abundance of each material present on the visible surface of the material at each point in time. From this prediction a probability score indicating the confidence of a significant abundance of material being present is computed. This information is then used to classify what type of satellite or other space object is observed. The proposed pipeline provides an explainable model that, even in the case of an unknown space object, would return useful information about its material composition. If the unknown object meets the expectations and

¹Mechanical and Aerospace Engineering, University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow G1 1XJ, UK. ²Electronic and Electrical Engineering, University of Strathclyde, James Weir Building, 75 Montrose Street, Glasgow G1 1XJ, UK. ³Fraunhofer Centre for Applied Photonics, Level 5 Technology and Innovation Centre, 99 George Street, Glasgow G1 1RD, UK. ✉email: massimiliano.vasile@strath.ac.uk

assumptions of a known class, it will be labelled as belonging to that class. Otherwise the proposed pipeline allows unknown objects that do not fit any known class well to be categorised as an Unidentified Flying Object (UFO).

This paper builds on the authors original ideas presented in^{19,20} and uses a mix of simulated data, based on object and sensor models developed in the previous work, and real measured spectra. The main contribution of this paper is twofold: (a) we improve the material identification approach previously presented in²⁰ including a comparison between a machine learning and a non-machine learning based identification approach, the ageing of the materials and a technique to account for unknown materials; (b) we propose a probabilistic approach to the identification of materials and classification of space objects.

We will show how, starting from high fidelity simulated data using real spectra acquired in lab environment, we can reliably unmix the spectra and identify the materials even in the presence of measurement errors, material ageing and atmospheric attenuation. Then we will show how objects can be classify with high probability from the expected combination of observable materials.

The paper is structured as follows. In “[Overview of the object identification and classification pipeline](#)” we provide an overview of the methodology, in “[Measured spectra, simulation models and data pre-processing](#)” we recall the way we generate and process the hyperspectral data, this section is largely based on the work done in¹⁹. In “[Material decomposition methods](#)” we present both the machine learning and the non-machine learning approaches to decompose the spectra and identify the materials. Finally in “[Probabilistic classification of space objects based on material identification](#)” we present the classification system.

Overview of the object identification and classification pipeline

An overview of the data processing pipeline presented in this paper will first be given in this section before presenting each step in more detail in later sections. A dataset of time-series spectral data of space-objects is required for this study. Ideally, this would come from real objects in orbit, but practically this presents several challenges. Firstly, in order to confirm that the models generalise well, the diversity of this dataset must be large. This means not only many different objects should be present in the dataset, but the training set must also contain a large range of orbital scenarios, trajectories relative to the observer, and illumination conditions even for the same object. Secondly, accurate evaluation necessitates ground truth data for those values which the models are to predict. In this context, this means having precise knowledge of the geometry, material distribution, and attitude motion at all times. Although the former two are theoretically possible to obtain, generally the materials used and their distributions over the surface of satellites is not something that is reported in literature. Additionally, if one wishes to include uncontrolled debris objects (eg spent rocket bodies) in the dataset, precise knowledge of the attitude will not be obtainable for real data. For these reasons it was decided to construct numerical physics simulation software to produce training data for the ML models.

The simulation models, as will be detailed in in “[Measured spectra, simulation models and data pre-processing](#)”, propagate the orbital and attitude motion of arbitrary space objects, and simulate the sensor output (aggregated onto a single pixel) at each timestep. This produces a time-varying spectral response as the object moves and rotates relative to the observer, and the illumination conditions change. The simulation model can simulate either a ground-based sensor (accounting for wavelength-dependent atmospheric extinction and Earth rotation), or a space-based sensor which follows its own trajectory.

The output of the sensor simulation is then decomposed into the proportion of signal due to each of the present materials. This can be done using machine learning or a more traditional constrained least squares approach, where both methods make use of a spectral library to determine the instantaneous material breakdown of the signal. The probability of a set of known materials being present on the object is then computed. The final step is to pass a detected combination of materials to a machine learning model trained to associate a given combination of materials to a class. This final step yields a probabilistic classification coming from the probability of detecting a combination of materials. A simple flow diagram of the proposed pipeline can be seen in Fig. 1.

This staged architecture was developed in favour of a direct classification to produce a final pipeline with greater transparency of the classification process. Not only does this pipeline enable a logical flow to be followed through the data processing, but also outputs useful information at the intermediate steps, such as the material abundances.

Measured spectra, simulation models and data pre-processing

Test data in this work was produced using a high-fidelity simulation model, developed by the authors in previous work^{19,20}. Since the focus of this paper is the data processing pipeline, only a brief overview of the simulation framework will be given in this Section; for more details the authors refer the reader to their previous works.

The model takes as its input a 3D model representation of a space object to be simulated along with an initial state describing the osculating orbital elements, attitude, instantaneous angular velocity and inertia tensor. The model then propagates attitude and orbital motion, and at each timestep calculates the illumination conditions of each facet of the 3D model. Before input to the simulation model, the facets are each assigned a particular




Figure 1. Classification pipeline and probabilistic material identification.

composition of materials, and a composite reflectance spectrum is constructed from the corresponding linear combination of materials. This, along with the instantaneous illumination conditions (illumination direction and view direction) are used to determine the visible and shadowed facets at each timestep, and aggregate the received light from all visible facets according to the Lambertian reflectance model. Reflectance and emission spectra are taken from lab experiments of known materials. Not only were these lab experiments used to gather the spectra required but also to investigate the consistency of the measured reflected spectra for different angles of incident light.

Spectra were acquired from laboratory experiments as follows. Materials were affixed to a cube placed on a rotation stage and many spectra were gathered while the cube was rotated to constantly change the relative angle between the face and the light source. While the total amount of light hitting the surface will inevitably change the overall amplitude of the gathered spectra we observed consistent spectra responses in terms of the location and relative magnitude of spectral features, as illustrated by the example for an aluminium foil shown in Fig. 2. This gives us confidence that the spectra we use within the model and the simulation method described above is consistent with a real world scenario. Where painted surfaces are included in our library we use reference spectra from the recognised United States Geological Study (USGS) library.

Figure 3 shows an example of three types of simulated space objects. The actual spectra coming from the laboratory experiments were placed on each of the triangles composing the surface of the object in Fig. 3. Light intensity derived from the simulated illumination conditions in space was then added to spectra on each triangle.

For the sake of the tests in this paper all the objects were placed in the same Low Earth Orbit and the sensor was placed on the surface of the Earth. Orbits of the imaged objects were randomized to increase training data diversity. The training set is generated for one fixed location of the telescope on the surface of the Earth, defined by its latitude and longitude. In deployment, images are likely to be taken during twilight hours to maximize the reflected light from the object while minimizing atmospheric scatter and sunlight leakage into the telescope aperture. Hence for all simulations the telescope was initialised on the western terminator, i.e. observing objects




Figure 2. Plot of multiple spectral measurements of the same aluminium foil at different orientations.




Figure 3. Example of three simulated satellites: (a) Iridium-NEXT, (b) Starlink, (c) DubaiSat-2. Colours correspond to regions with different materials. The reflectivity spectra for each region were obtained from laboratory experiments using mock-ups.

in the early morning, coming over the horizon towards the sun. The orbital elements of the objects are generated randomly within the bounds in Table 1, which ensures that the objects begin their propagation between 25° and 10° before zenith with respect to the observer. 5-min observation arcs are simulated for all training samples, however the pipeline proposed is intentionally agnostic to the duration of the exposure and number of time samples.

The initial orientation and angular velocity vectors are randomly initialised, with an upper bound on the rotation rate of two rotations per minute. Although many of the satellites simulated in this paper would in reality have nadir-pointing configurations, one possible application of this technology is locating lost satellites with which communication has been lost, and these targets may be tumbling. Taken together, the randomisation of all these parameters ensures a good level of diversity in the training set. For all simulations, an observation arc of 5 min was simulated, sampled at 1 Hz. The band width of the hyperspectral sensor was set to 5 nm. The simulation model is built to be compatible with arbitrary wavelength bands, but in these simulations we used the spectral range spanning 455 nm to 1035 nm, with a constant 5 nm band width. The incoming spectrum is simply integrated across each band's range to obtain the total collected power. This is then converted to sensor counts by dividing by the band's central photon energy. Shot noise and readout noise are accounted for in the sensor model.

The hyperspectral sensor was modelled as yielding no spatial information, thus in all experiments in this paper the light collected by the simulated sensor in a given wavelength band is captured by a single pixel. Thus, the received signal is two-dimensional, varying only in wavelength and time, as opposed to the more typical definition of hyperspectral imaging where there is a three-dimensional signal of two spatial and one spectral dimension.

A key pre-processing step for the simulated sensor output is the transformation from band-wise photon counts to a hyperspectral colour-indexed representation of the signal, which is an application of colour indexing to hyperspectral signals. This transformation removes distance and absolute brightness-related effects while preserving the spectral shape and relationships in brightness between different bands, which is the quantity of interest to this work. This transformation acts analogously to the normalising or scaling transformations that are typically applied to data before being passed to machine learning models for training or testing. The equation for the transformation is:

$$\mathbf{s}(t, \lambda) \rightarrow -2.5 \log_{10} \left(\frac{\mathbf{s}(t, \lambda)}{\mathbf{s}(t, \lambda_{ref})} \right) \quad (1)$$

where $\mathbf{s}(t, \lambda)$ is the time-evolving spectrum received from the object, and λ_{ref} is an arbitrary reference wavelength band. For the simulation of the sensor on ground we include also the effect of the atmosphere.

With this model it is possible to define the attitude of the space objects for which the observation is simulated¹⁹. As the purpose of the proposed models is to predict the materials present from spectral data and what type of objects this may correspond to, we elect to use randomised attitude for all objects rather than attempt to mimic the real world pointing directions of particular satellites. This was done to improve the robustness of our study as a greater number of observation scenarios could be investigated. This allows us to confirm if the proposed models can robustly identify materials on all faces of all objects simulated, not just those visible from a defined perspective.

Material decomposition methods

In order to reconstruct the material-wise contribution to the received signal over time, two methods were investigated: one using machine learning and one using a more traditional, explainable method. These two methods are explained in this section and results of each will be presented. This step of the pipeline will allow us to extract the material abundance curves (MAC) associated to each observed object, or curves providing, for each observation time step, the relative abundance of each material present in the received spectrum.

Machine learning material identification

We first approached the identification of individual materials in a spectrum by training an Artificial Neural Network (ANN) to decompose the spectrum. In order to train an ANN to decompose a composite spectrum into its material components, a training set was constructed composed of N random linear combinations of m

Parameter (symbol)	Value range
Semi-major axis (a)	[7471, 7671] km
Eccentricity (e)	[0, 0.01]
Inclination (i)	[−45°, 45°]
Argument of Periapsis (ω)	[0°, 360°]
RAAN (Ω)	[−5°, 5°]
True anomaly	360° - ω - [10°, 25°]
Telescope latitude	[−5, 5]
Telescope longitude	[−5, 5]

Table 1. Range of initial orbital elements used for training data generation. The true anomaly first places the satellite at the ascending node $\theta = 360^\circ - \omega$, and then moves it backward in its orbit by 10° to 25°.

materials, constrained such that their sum is equal to 1. While it is an option to train the decomposition ANN using simulated sensor data from the simulation model, it was found that this approach effectively embeds the geometry of the underlying objects in the training set, reducing the generalisation performance when the model is used with unseen geometries. When instead training on these linear combinations of raw real spectra from laboratory experiments, a geometry-agnostic model is produced which performs similarly well across all geometries we exposed it to for validation.

The ANN developed for this stage had a relatively simple architecture—a fully-connected, feed-forward regression network, with two hidden layers and a combination of multiple activation functions on those hidden layers. The softmax activation function was applied to the output layer to ensure that the sum of all material contributions was equal to 1. The full list of hyperparameters in the final model can be seen in Table 2. The model was trained with early stopping enabled such that the weights and biases revert to the best found set after eight epochs of no improvement. The final model developed and used for later stages of the data processing was trained with a library of nine materials: black, white, red and green paints, aluminium foil, Multi Layer Insulation (MLI) (or gold colour) thermal blanket, copper and titanium metals, and GaAs solar panel. Paint reflectance spectra were obtained from the USGS spectral library²¹, while the remaining spectra were obtained from our own laboratory measurements.

After training on the dataset of spectral combinations, the ANN was capable of making accurate predictions of the material-wise breakdown of the received light for all tested geometries. Example curves for a simulation of an Iridium-NEXT satellite can be seen in Fig. 4, showing a good fit to the ground truth with the exception of black paint, which was found to be generally overestimated, although, the overall shape of the curves is correctly predicted. Additionally, albeit the library typically contains more materials than are present in an individual test satellite, it did not predict the presence of materials that were not in the satellite being tested (for example, if a satellite does not contain white paint, the ANN did not predict any significant quantity of white paint despite it being in the library).

Non-machine learning material identification

While good unmixing performance was achieved using machine learning techniques such methods often lack explainability and require time and data to re-train. Non-machine learning approaches were explored to ascertain if similar unmixing could be achieved in a more explainable way and improve confidence in the ability to unmix material abundances reliably from spectral data. A number of techniques exist in the literature for the unmixing of hyperspectral data. To successfully unmix a received signal two problems need to be solved; finding the fundamental spectra present in the data, i.e. the spectra of the materials that are present, and what proportion of these spectra make up the received spectra. The fundamental spectra that make up the received signals are more commonly referred to as endmember spectra and a wide range of techniques focus specifically on identifying these spectra. While the specific methods differ, these techniques typically assume that the endmember spectra will exist in relatively pure form (i.e. the spectrum is composed of near 100% one material) in the data at some point in space or time. For example, well established solutions such as N-findR²² are very highly performance limited by the purity of received spectra. This technique assumes the volume of a simplex formed by the endmembers is larger than the volume of a simplex formed by any other combination of spectra. Using this assumption, the algorithm can search combination of spectra in the received data to find those that best represent the endmembers of the data. However, this means all returned endmembers are selected from the data, thus this method will only represent the purest spectra that exists in the data which may still be some combination of multiple materials. Methods other than the volume of a simplex can be used to compute the likelihood of a spectra being an endmember, such as orthogonal subspace projection²³ or vertex component analysis²⁴ among others²⁵, but the same limit will apply while endmembers are selected from the received spectra. Alternative approaches exist using methods such as SVD²⁶ or Tucker Decomposition²⁷, which have recently been applied successfully to unmix materials in hyperspectral images of scaled down satellites in a laboratory environment. However, in²⁷ the authors note that the higher the proportion of pure pixels the better the more accurate the resulting endmember spectra will be. The test data used in the worst case still has 25% unmixed pixels in the data. Such methods are thus not best suited for single pixel spectral measurements where the level of mixing

Hyperparameter	Value
# hidden layers	2
# hidden neurons	200, 20
Activation (h1)	ReLU
Activation (h2)	Sigmoid
Activation (output)	Softmax
Optimiser	Adam
Loss function	MSE
Train set size	8,100,000 samples
Test set size	900,000 samples
Minibatch size	100,000 samples

Table 2. Hyperparameters of the decomposition ANN. The number of input and output neurons must be equal to the number of spectral bands and library materials respectively.




Figure 4. Example of ANN predictions for material abundances for an Iridium-NEXT satellite. Over 200 simulations with differing orbital and rotational conditions, the mean squared error in the predictions was 5.32×10^{-3} .

is likely to always be high. To resolve these challenges, it is proposed to unmix the received spectra based on a library of candidate materials that are likely to exist in space. This approach is conceptually similar to the ANN proposed in “[Machine learning material identification](#)”. The idea is to compute, for each point in the time-series, the weighted linear combination of reference spectra in the library that would produce the received spectra. If successful this returns the abundance of each material present in each received spectra, with materials that are not present returning a 0% contribution. A schematic of the proposed library unmixing method is shown in Fig. 5.

As with the machine learning models it is necessary to first remove any atmospheric and solar illumination effects to achieve good performance although it is assumed this occurs prior to implementing the method in Fig. 5. Band selection is used to select and use only wavelength bands for which information exists in the received spectra and all spectra in the library. In addition, the non-machine learning methods deploy data smoothing and continuum removal techniques to ensure that unmixing is based on key features. A detailed description of these techniques, with illustrating figures, can be found in¹⁹. Data smoothing will remove high frequency noise while continuum removal, achieved using a convex hull transform, will remove the overall concave shape often present in spectra while normalising the signal. This must be applied to both the library spectra and the received signal. This was not necessary in the machine learning technique, as the network learns which features are important during training. An example of spectra before and after this processing can be seen in Fig. 6. Note

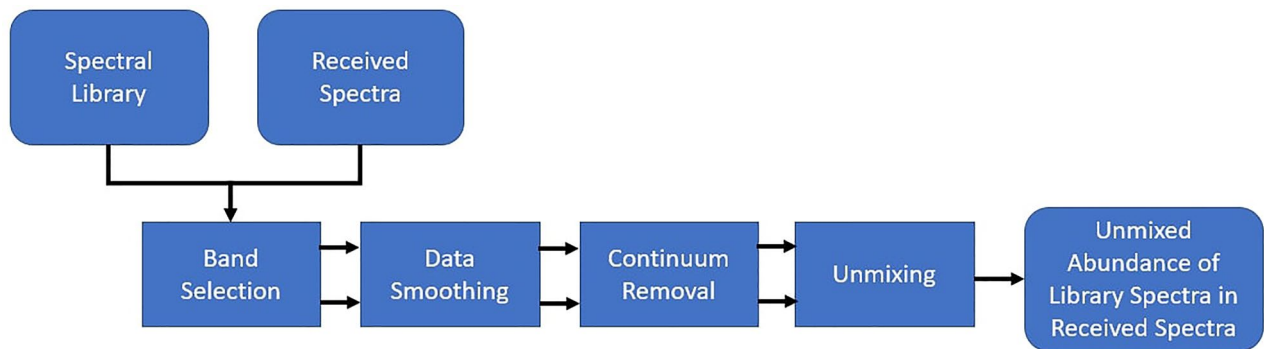


Figure 5. Schematic of library unmixing algorithm.




Figure 6. Example of pre-processing where (a) original aluminium spectra and (b) is spectra after all pre-processing steps.

that the overall signal has been flattened, high frequency noise removed and thus the key absorbance band at 550 nm is more pronounced.

The unmixing of the spectra based on the library can theoretically be performed using multiple unmixing techniques. In this work, methods based on non-negative constrained least squares (NCLS) and Joint-Sparsity and Total Variation (JSTV)²⁸ were evaluated, and results were found to be very similar. The unmixing approach was validated on a small dataset of simulated data. An example unmixing result based on NCLS is shown in Fig. 7.

The results show good accuracy with minimal or no estimated contribution from materials that are not present in the data. The library unmixing is also able to track changes in the abundance of materials that are present as the satellite rotates well, giving results that appear similar to those in the machine learning method. A full comparison between these methods is provided in “[Comparison between machine learning and non-machine learning unmixing](#)”. To decompose the spectra into abundances of materials, the library unmixing must find the optimal weighted combinations of materials to reconstruct the received signal and thus computational complexity scales with the number of materials in the library. In the case of the ANN method these combinations are checked in advance during training so the computational cost comes at that stage rather than during deployment, when the model most simply map the received signal to closest learnt combination. For the least squares based library unmixing the increase in complexity from larger libraries can significantly change execution time of the software. For example the nine material library used in this study would take approximately 3 ms to unmix each spectra, while increasing the library eightfold to 72 materials would increase computation time 50-fold to 140 ms.

Unmixing with incomplete spectral libraries or unknown materials

A limitation of library unmixing methods is the requirement to have an accurate spectrum in the library for each material that may be present or a large enough library to contain all possible variations of the spectra for each expected material. Given that the scope of the unmixing is to identify materials typically used in the space sector, building a sufficiently rich library with all representative materials is possible.

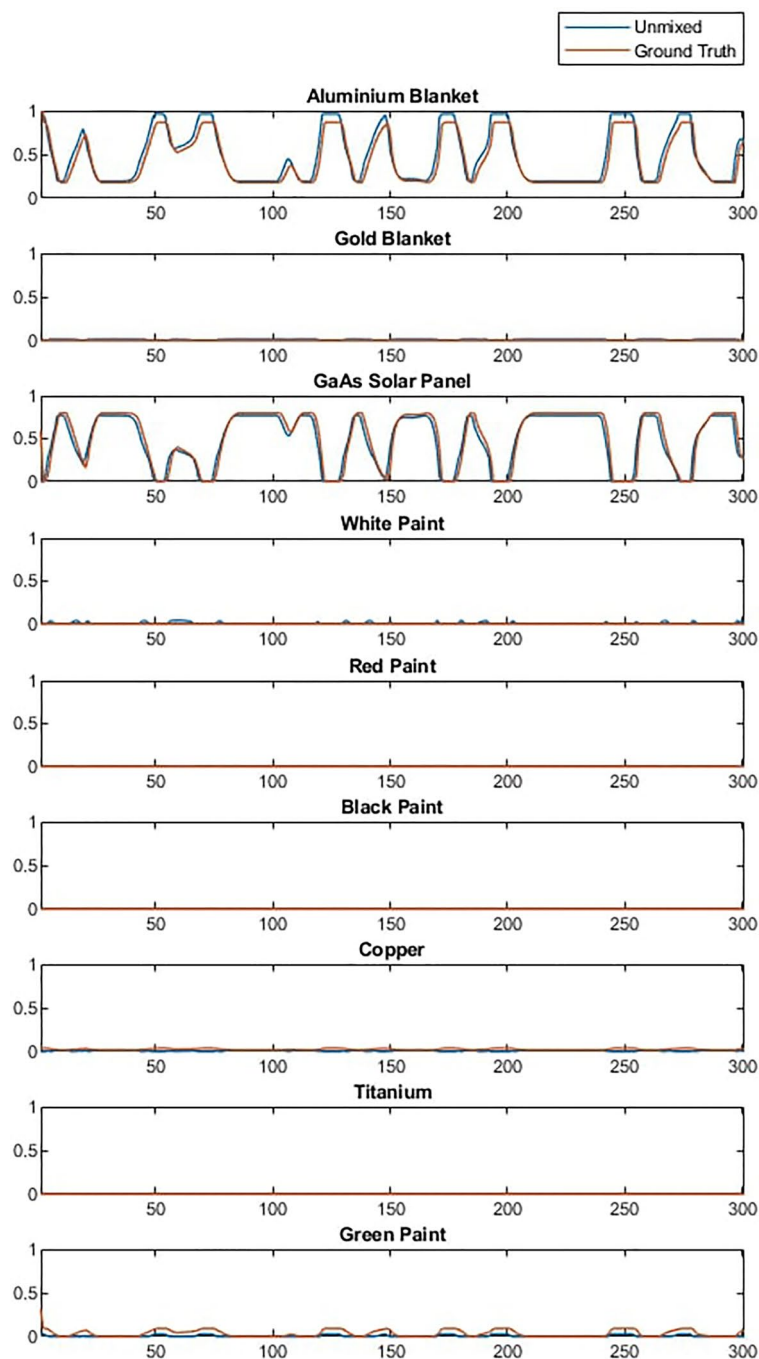


Figure 7. Example of unmixing a received timeseries signal.

However, if a space object contained materials that were not present in the library it would not be possible to identify those materials. Furthermore, as the objective of the unmixing algorithm is to recreate the received signal as accurately as possible using weighted linear combinations for the library spectra, the algorithm would likely adjust contributions of other materials in an attempt to recreate the features of the missing ones. It is also possible that the aging of materials in space could result in differences in the spectra that are included in the library. In this section a preliminary exploration of both challenges is presented alongside a potential solution through the analysis of a residual portions of the signal that cannot be reconstructed.

The effect of materials that were not included in the spectral library can be studied by removing spectra from the library and using the algorithm described above to unmix the data described in “[Measured spectra, simulation models and data pre-processing](#)”. Figure 8 shows the result of unmixing a Starlink satellite with and without the solar panel in the spectral library. It can be observed that with solar panel present the abundance estimation




Figure 8. Example of unmixing a received timeseries signal of Starlink where (a) is unmixed with a complete spectral library and (b) is unmixed with GaAs solar panel removed.

is good but that when the solar panel is not present in the library the algorithm attempts to compensate for this by mixing in a combination of white and black paint.

The reason for this can be explained by studying the received spectra after processing (blue) in Fig. 9. The received spectra has features which the unmixing model recreates as accurately as possible by using combinations of spectra in the library. Combining white and black paint allows the model to recreate the overall shape of the spectra, but misses fine features that are unique to the solar panel.

This behaviour is problematic and highlights the importance of having as complete a spectral library as possible when using this technique. However, this result also highlights potential solutions to mitigate this issue. Due to the missing material the mixed spectra did not match the input spectra well, thus the mean squared error between the received and mixed signal could indicate some cases where errors occur and that the result should be studied in more detail. As key features of the received spectra were not recreated in the mixed spectra the residual spectra (i.e. the difference between the mixed and received spectra) would capture these features. Features in the residual spectra will be caused by noise, errors in unmixing and missing spectra from the library. When features result from the presence of a material that is not in the library it would be expected that the residual spectra would resemble the spectra of the missing or unknown material. This concept is demonstrated in Fig. 10, which shows a comparison between the residual spectra in our Starlink example and the spectra of the solar panel known to be in the received spectra but not in the library. Note that as residual spectra are small and can have positive or negative values the spectra have been scaled to allow comparison of the features present. It is clear in this example that there is significant overlap between the residual features and the solar panel spectra, although these are not identical, likely due to the presence of noise and small unmixing errors.

It is important to note that in such cases the spectral library is incomplete and the primary use of this residual spectra should be to facilitate further analysis, understand what spectra might be missing and update the model.

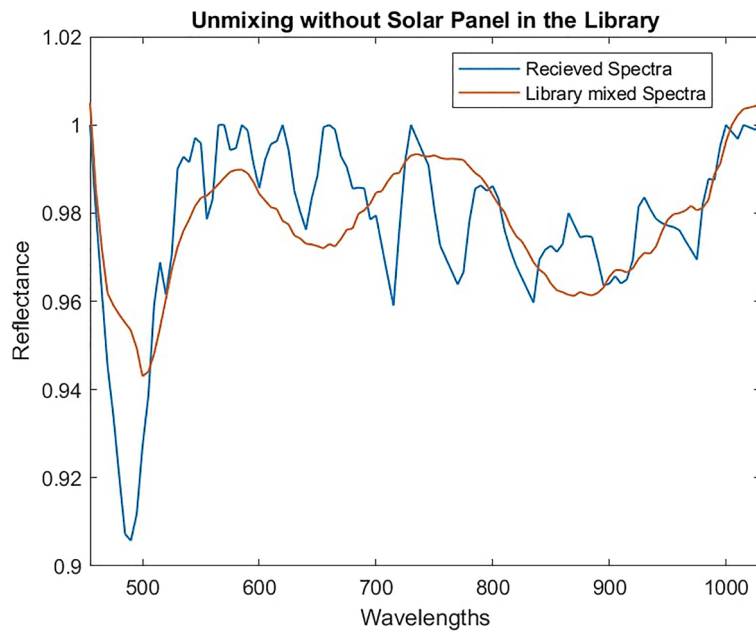


Figure 9. Accuracy of unmixing using white and black paint to compensate for missing solar panel.




Figure 10. Analysis of residual spectra compared to real solar panel spectra that was missing (scaled to show similarity in features).

However, there is potential to use the residual spectra as a short term fix for the overfitting of white and black paint in this case. To do this the unmixing model is applied twice. On the first pass unmixing is based only on the spectra library and the residual spectra computed. On the second pass the residual spectra is normalised and included in the spectral library, thus allowing some abundance to be assigned to this rather than a material that is not present. The result of applying this residual based correction method to the Starlink unmixing is shown in Fig. 11.

These preliminary results suggest good performance in preventing this overfitting problem by using the residual to correct the initial unmixing. However, the performance of this method may depend significantly on which signal was missing, or indeed how many signals were missing. In this case, the missing solar panel was a very significant material and thus the residual spectra matched its feature well. If instead the missing material




Figure 11. Library unmixing of Starlink satellite with residual correction.

made up only 1–2% of object then this method may not help identify this material as well, as even 1% of the original spectra resulting from noise, or a 1% error in mixing, would create a residual that depends as much on these random processes as the missing material.

Comparison between machine learning and non-machine learning unmixing

To compare and analyse the differences between ANN and Non-machine learning library matching techniques a new dataset was generated with a few examples of various object types. Results from both methods on a larger dataset is also shown in “Probabilistic classification of space objects based on material identification”. Both methods were used to unmix the spectra in the timeseries and the resulting abundance estimations were compared to the known ground truth. Results are broadly similar which is confirmed by the similarly low mean squared errors shown in Fig. 12 and the similar abundance graphs in Fig. 13, which shows examples of unmixing results for three different objects. The mean squared error of the ANN was 5.3×10^{-3} , nearly exactly the same as the

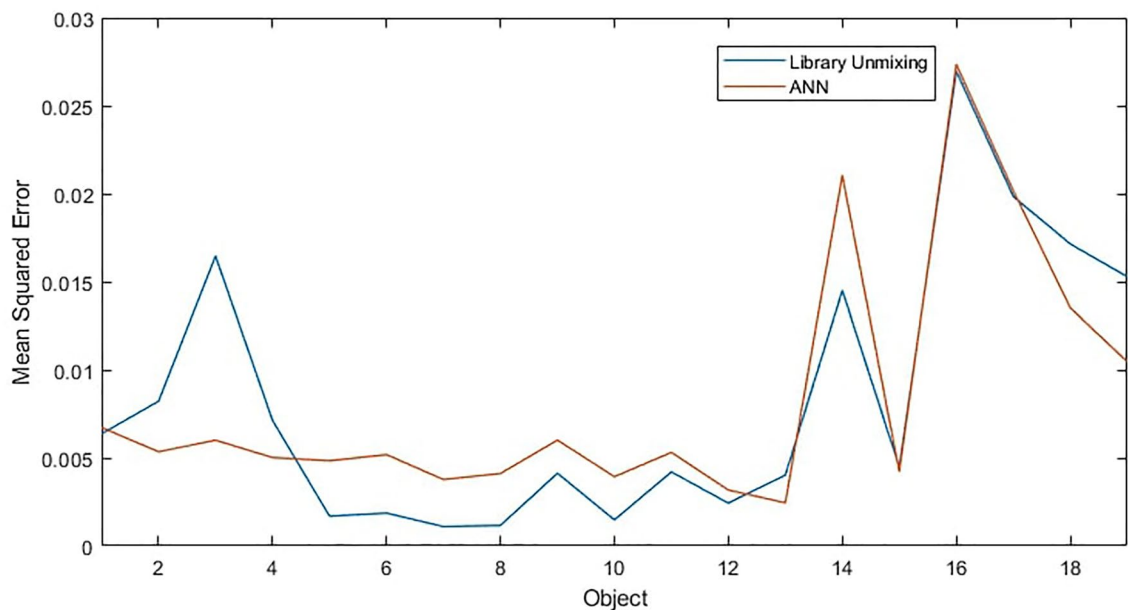


Figure 12. Mean squared error of ANN and library matching unmixing approaches for different simulated objects.

previous trial, and the mean squared error of library unmixing predictions was 4.5×10^{-3} . This is a small difference and Fig. 13 shows the the best unmixing result differs from object to object, suggesting either model could be used to obtain similar results and which method has the lowest error would vary between datasets. However, while the overall magnitude of errors are similar the type of error encountered by both methods are quite different. Figure 13 (a) shows a typical example were differences between the methods are very subtle while (b) shows the case where the ANN performed slightly better than the library matching and (c) shows the case where the ANN performed slightly worse. It can be observed that when the ANN produces the errors it typically does so by predicting a small amount of a material that is not there (such as white paint), while the errors in library matching more often come from over- or under- estimating particular abundances, or missing a material that exists in a smaller quantity. This suggests library matching may produce larger errors, but less frequently compared to the ANN. Overall the performance of both methods is found to be accurate and reliable for further classification and analysis, however, these differences in characteristics are worth considering in future developments. For computational complexity, the non- machine learning library unmixing takes around 3 ms to analyse each spectra compared to 0.7 ms for the ANN method. This performance advantage for the ANN is likely due to more of computation taking place at the training stage of the ANN, thus this advantage will increase further as the complexity of the problem increases, such as if larger material libraries were used.

Material aging and inaccurate library spectra

Materials in space exhibit spectral changes over time due to exposure to the space environment^{29,30}. Although the ANN could easily be retrained with a new set of spectra corresponding to a given exposure duration, if the object being imaged is not already well-known, it is unlikely that good knowledge of its age will be available. Additionally, to the best of the authors' knowledge, no models currently exist to describe the spectral aging process and predict the spectral changes that will occur over time. Thus there is likely to be some error in the library spectra with respect to the underlying materials on the real object. It is, therefore, important to understand the effects of such error.

Since no model is available, as a first-order approximation of the spectral changes we define an aging function $A(\lambda)$, which is used to transform the reflectance spectra $R(\lambda)$ of all materials in the library: $R(\lambda) \rightarrow R(\lambda)A(\lambda)$. The function was chosen to accentuate reflectance in the red or near infrared spectral regions relative to shorter wavelengths given that in the literature there is the evidence that the reflectance spectra of aged materials display a shift to the red. An aging function of sigmoid shape is defined such that an increasing boost is given to the spectrum in the red and infrared regions:

$$A(\lambda) = 1 + \frac{a}{1 + e^{-0.02(x-750)}} \quad (2)$$

We reiterate that this is simply a first order approximation of the spectral reddening effect reported in some literature on object spectrometry and is not to be taken as an accurate physical model of the process. The aging function along with a comparison of aged and unaged aluminium blanket spectra are illustrated in Figure 14.

To investigate the effect of this spectral error on the ANN decomposition, we generated several datasets of simulated measurements for DubaiSat-2, where the underlying materials in the simulation model have been replaced by increasingly aged versions ($a = 0, 0.025, 0.05, 0.075, 0.1$). The existing ANN (trained on unaged

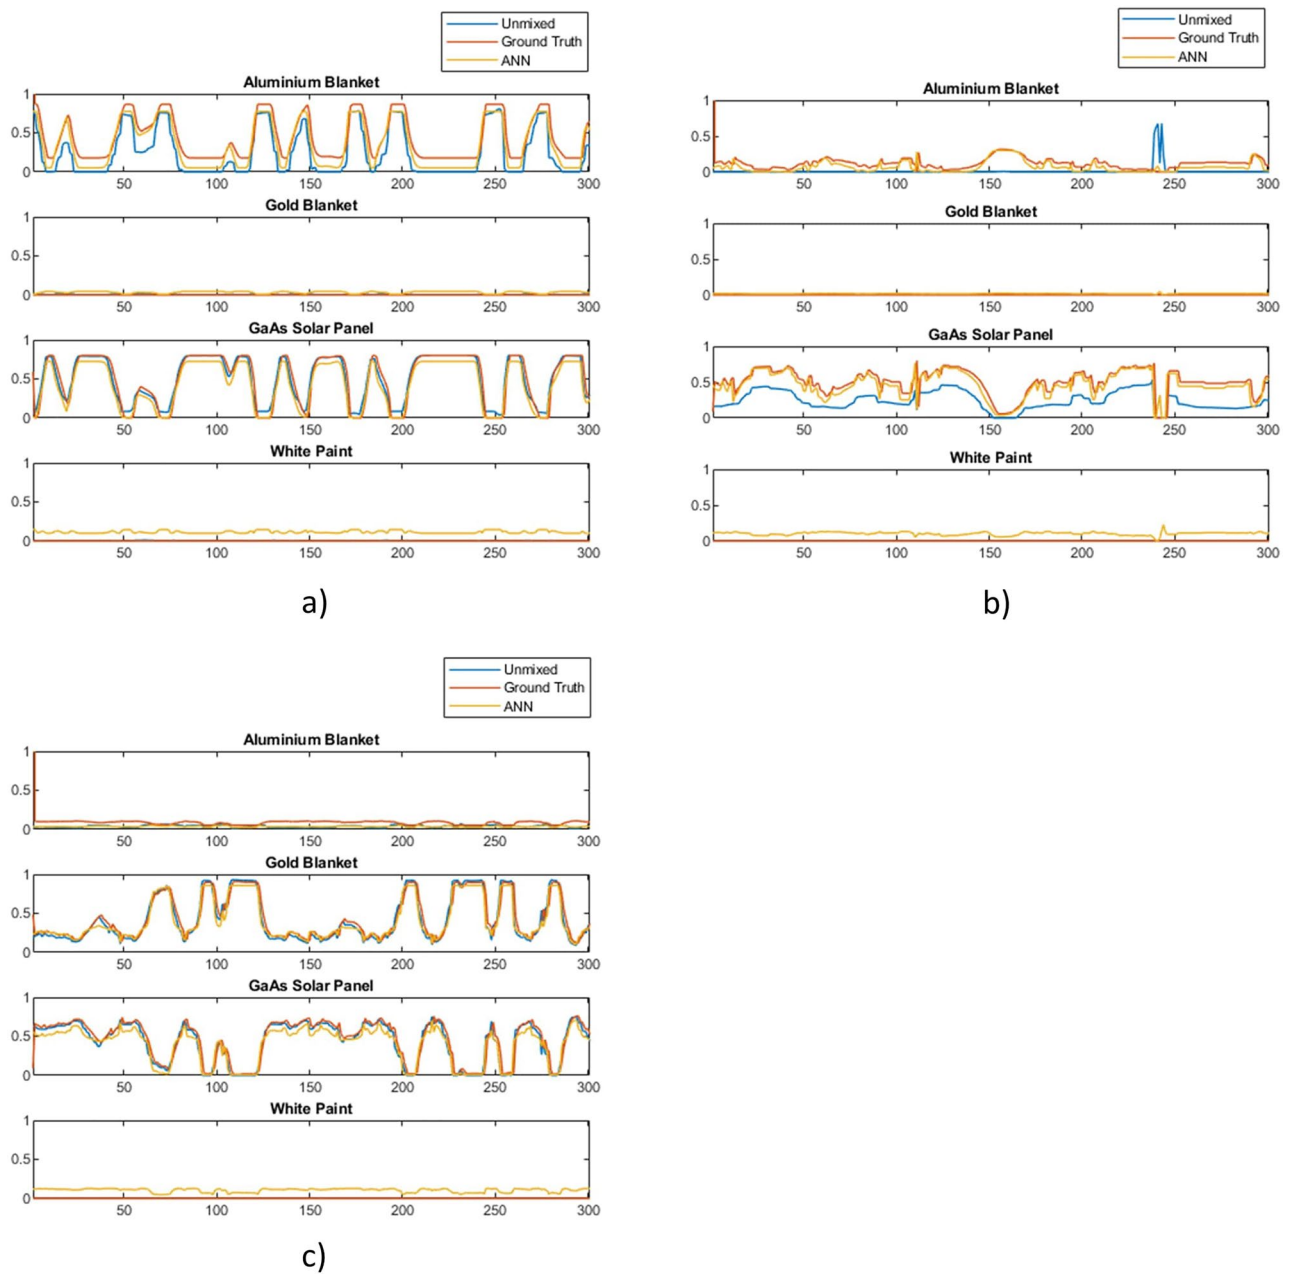


Figure 13. Unmixing performance for aluminium, gold and solar panel for 3 objects with ANN and library matching methods where (a) similar performance, (b) ANN slightly better and (c) ANN slightly worse.

spectra) was then used to decompose the signal into its material components. Model predictions for these four aged versions of DubaiSat-2 are shown in Fig. 15.

These results show that discrepancies between the spectra used to train the MAC model and the spectra of the materials in a real object do cause errors in the predictions made by the MAC model. Of particular note is the progressively increasing overestimation of the abundance of black paint, as the aging is increased. Since the sum of the material decomposition is constrained to 1, this causes other materials' contributions to be reduced to compensate. However, it should be noted that the overall shape of the decomposed material curves is preserved.

Repeating this test with the library matching method shows similar results with typical examples of the error shown in Fig. 16a. Note that a different instance from our simulation, but of the same object and aging effects, is shown than in Fig. 15, as the chosen example best illustrates the typical errors observed when using the library matching method. Even with the greatest amount of aging in our simulation ($a = 0.01$), it is observed that the overall estimate of material abundance is accurate, but changes in the abundance at some time points are missed. This is particularly noticeable in the aluminium. The relative performance difference here is likely explained by the pre-processing used, which includes data smoothing and continuum removal, which will reduce some of aging effects. However, errors still exist and is likely caused by aging creating more similarities between the spectra than before. As with the ANN method it is possible to re-train with knowledge of the aged spectra to

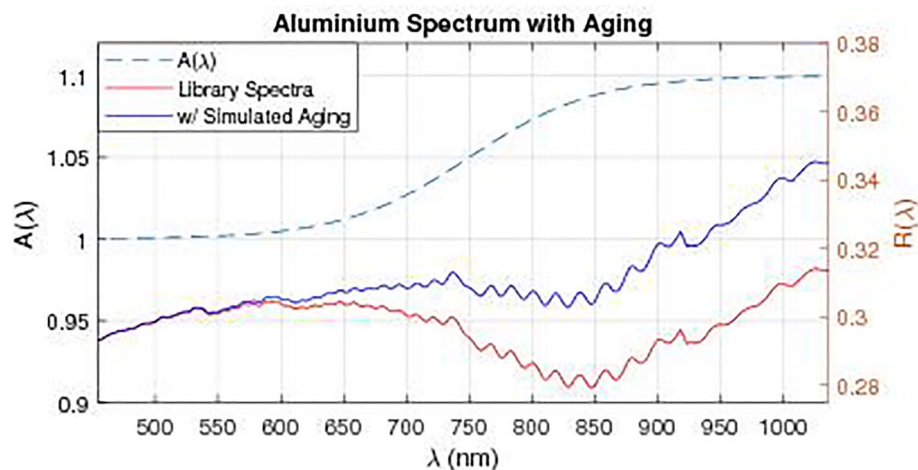


Figure 14. Aging function (with $a = 0.1$) and its application to the spectrum of aluminium blanket, in the wavelength region of interest for this paper (455–1035 nm).

obtain better results as shown in Fig. 16b. In this case that is achieved by adding aged spectra to the library in place of the original spectra. This demonstrates that our unmixing model is relatively robust to the effects of aging, but the best performance will likely be achieved if the aging of materials is well understood and the model can be re-trained, or spectral library updated, accordingly.

Test on known satellite components

We tested the material classification methodology on spectra measured from a known satellite component. The component is an AzureSpace triple junction GaAs solar panel shown in Fig. 18. We imaged this component both in a controlled laboratory environment and also in outdoor conditions with a telescope capable of imaging space objects. This allowed us to explore the robustness of the model to the noise induced by real sensors, lenses and solar illumination while having good ground truth data and a precise understanding of what we are unmixing. For the laboratory imaging, the spectra of this panel were captured using a Headwall VNIR hyperspectral camera with a range of 400–1000 nm which captured spectra per pixel, with each pixel representing a spatial region of ~ 0.1 mm. This allowed thousands of spectra to be captured across the surface of this panel so that the contribution for the GaAs panel and the connector material could be analysed independently. The mean spectra of solar panel and connector regions, without any processing, are shown in Fig. 19a and these spectra were added to the library. The GaAs spectrum captured is quite different from the spectra used in the existing library which had evidence of small amounts of reflections at multiple wavelengths. In this example we see very little reflection other than some reflectance at 900nm. This suggests that higher quality solar panels such as this have less reflectance, and thus may better represent the panels used on satellites. From Fig. 19a it is also clear that the connector spectra is significantly different from that of the GaAs and that the peak of the connector signal is significantly greater than the peak of the GaAs spectra. This suggests that it is possible that when measuring a large solar panel array with many connectors that the spectra captured could be dominated by the connector material, rather than the GaAs material. In this case it may be possible to use the detection of the connector material as an indicator of solar panels. For imaging using a telescope, Fig. 17 shows the reflectivity spectrum of the solar array measured with the Celestron HD 11" telescope based spectrometer system at the distance of ~ 25 m under solar light illumination in clear sky conditions. For this spectral measurement, a lens was placed after the "image plane" of the telescope, which transferred the image of an object to the fibre entrance of the commercial off-the-shelf spectrometer Ocean insight "QE-pro". The spectrometer measured a spectrum of sun light reflected from an object. Reflectivity spectrum of this object was calculated as the ratio of the measured reflected sun light spectrum to the spectrum of the sun measured at daylight conditions. Several spectra of the same panel were gathered in this study. The library unmixing models proposed in this paper used a range of pre-processing techniques, described in "Non-machine learning material identification". In addition to this, for this experiment the spectral resolution is also sampled to match that of the spectra previous used in the library, to ensure a good match between features of the same material. Figure 19b shows a measured spectrum from the telescope after all pre-processing and the one obtained with the library matching approach presented in previous sections where other panel spectra, gathered by this telescope, is also included in the library. While the signals do not match precisely due to the effects of noise (we used spectra from single time points for this experiment), it is clear that the features themselves are all recreated. This shows that the algorithm performs as expected when the unmixing receives a signal that contains noise but otherwise has spectra of a material which exists in the library. However, while this demonstrates the ability to unmix and re-create the received spectra in the presence of noise it is quite idealistic as it shows an example where the library contains spectra from the same instrument and day as the target to be unmixed. A more difficult problem is thus also posed where no telescope data is included in the library and instead the connector and solar panel spectra from Fig. 19a are used in the library. 4 spectra were gathered using the solar panel using the telescope and the abundance of each material identified




Figure 15. MAC model predictions compared with true values for DubaiSat-2 with progressively increasing spectral aging. Only present materials shown. Different simulations of the same satellite shown in each subfigure.

is reported in Fig. 20. The results show a high abundance of the solar panel component. This suggests that, even without any prior knowledge of the noise incorporated within the training or library, the model is able to detect known materials. Small amounts of aluminium and green paint are also estimated which are likely causing by the model attempting to unmix noise, as explained in “Material ageing and inaccurate library spectra”, and genuine detection of other elements of the panel such as screw holes.

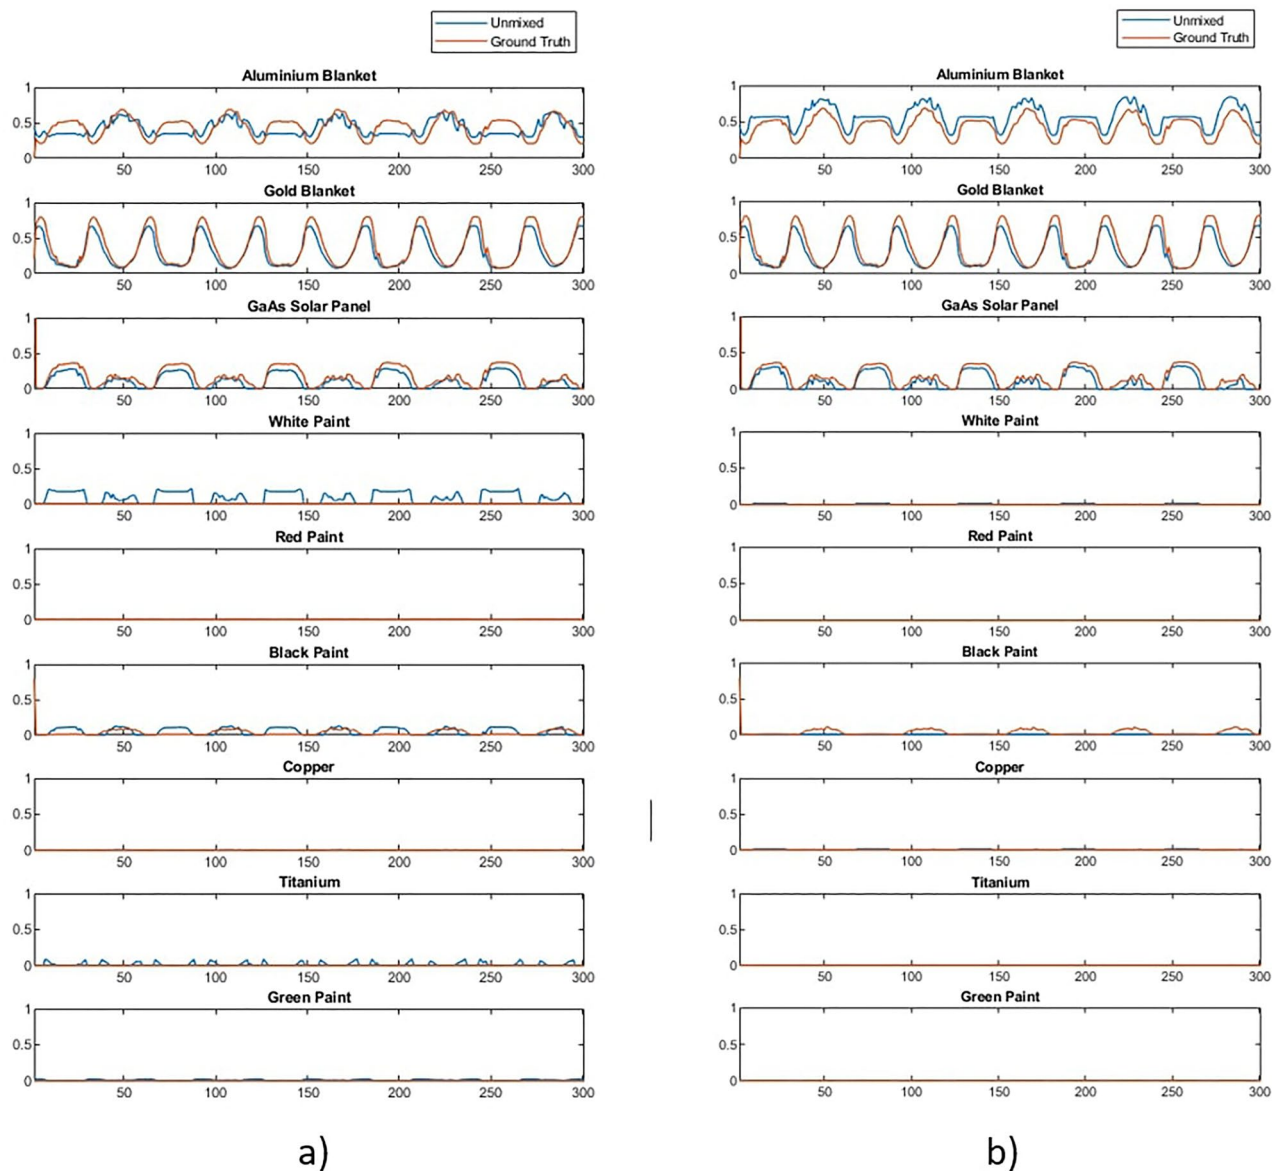


Figure 16. Unmixing of aged DubaiSat ($a = 0.1$) with (a) the original spectral library, (b) a spectra library that has also been aged.

Probabilistic classification of space objects based on material identification

Once a model is developed to extract the material abundance curves (MACs), its output can be further processed to extract information on the underlying space object. Specifically, we propose that knowing the abundance of material on the surface of a space object can allow the class of that object to be predicted, which in some cases may inform the purpose of the object itself. In theory, space objects with a similar general purpose are likely to have similar visible components, and thus similar materials. For example, satellites are expected to have solar panels while rocket bodies would not and thus the presence of any material from which solar panels might be made will help predict which of these we may observe. Similarly, those satellites deployed for communication purposes would be expected to have antennae, so materials from which an antenna might be made, could be used to help predict which satellites fit this classification. It is desirable to design a classification system that is agnostic to attitude and pointing direction, as the resulting model would then be applicable to multiple deployment scenarios. Thus the model used to simulate satellites will randomly generate its motion and not assume the same repeated behaviour. While the abundance of materials will very likely help distinguish objects, the former assumptions means that it cannot be guaranteed that the surfaces of the satellite that are observed will accurately reflect the true abundance. Indeed in any practical scenario this is also unlikely to be the case. For this reason, the classification model is designed to classify objects based on the probability of detecting materials.

To achieve this, we first demonstrate that the presence of materials on space objects can be reliably identified using the MAC developed in previous sections. A probabilistic approach is used in which the time-series data of material abundance is mapped to a single probability score indicating the confidence that the material is present on the object, rather than a small false positive abundance due to noise. This probability estimate for each material




Figure 17. Satellite component: GaAs solar panel measurement from roof mounted telescope.




Figure 18. Satellite component: GaAs solar panel with connectors.

is then used to classify the space object. While it would be possible to train on known object data from simulations, this has the potential to create over-fitting problems and may lead to a model that does not generalise well. Instead the proposed classification model is trained on synthetic instances representing the potential material probabilities of each class. It is also possible that a machine learning model trained on simulated objects might learn to identify each object and then map this to the corresponding task, which is the opposite of the intended workflow where we intend to perform a more general classification to progressively reconstruct a full picture of the observed object. Thus the model we train has never seen any simulated or real object data until the moment it is tested, as all training data are only synthetic mixtures of materials.

To be noted that the composition of the space objects used in the simulations in this section does not necessarily correspond to the actual composition of those objects. Shapes, attitudes and material distributions were chosen to be sufficiently representative of existing space objects but do not correspond to the truth. It is, therefore, possible that the objects listed in this section are not really covered in the materials we selected nor the abundance of those materials corresponds to the actual one. On the other hand the examples in this section are useful to demonstrate the ability of the proposed methodology to identify the materials and classify the corresponding object because we can compare the outcome of the proposed pipeline against the ground truth of the objects we simulated.

Probabilistic material classification

The first stage in this approach is to analyse the material abundance estimation, to gain a probabilistic estimate of what materials are present. The material abundance models proposed in this work do not inherently provide a




Figure 19. Spectra: (a) individual measured spectra of cells and connectors, (b) mixed library spectra vs ground truth.

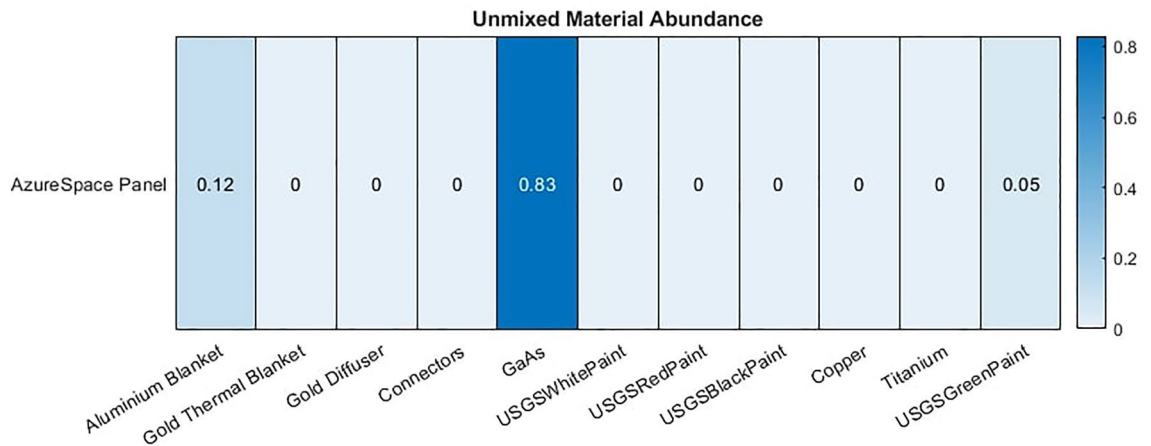


Figure 20. Satellite component: unmixing of telescope data of solar panel based on library containing laboratory spectra of elements of that component.

binary material presence decision. This is because these are unmixing models that attempt to find the proportions of each material from a known library that in a linear combination would produce the received spectra. Thus the probability of a material being present is linked to the abundance of the material that is estimated and must be extracted from this for our purposes. As it is known that the visibility of any material will alter as the object rotates the maximum abundance of each material detected for any point in time offers the best estimate of being able to conclude if that material is present. The more significant the features of a given material spectra are in the received spectra the more confident the model is that this material is present and subsequently a greater abundance is predicted. However, it is not helpful to consider 100% abundance as the threshold for 100% probability as most materials on most objects never reach this abundance for any angle of observation. Instead a far lower abundance should be sufficient to provide confidence that a material is present. In the ideal case the abundance that equates to 100% probability of prediction should be the lowest abundance that could not be generated by noise. We propose a threshold where the model has maximum confidence that a material exists and linearly scale the output of the material abundance prediction model accordingly. This allows separate measurements of probability of detection and material abundance, despite being derived from the same data. For example, with the threshold set to 5% and two satellites had 10% and 50% solar panel respectively, the model would return a 100% probability of the panel existing in both cases, but retain the knowledge that there is a larger panel in the second case. Empirical studies of the error rate of different values suggest a threshold of 5% appears effective in this study as it is consistent with the smallest abundances observed and greater than the typical error rate from the material abundance models. This would mean that the 1–2% false positive abundance predictions discussed previously would equate to a relatively low probability of predicting that material is present.

To investigate these probability estimates, a large dataset of 26,000 space objects were generated, 13 object types with 2000 simulated instances of each. From this simulation we record the abundance of each material that

is visible to the observed and used by the model to generate the spectra. The mean abundance of each material for each type of object is shown in Fig. 21. All objects were observed for 5 min at 1 Hz with the simulation using the same observation and rotational conditions described in “Measured spectra, simulation models and data pre-processing” and Table 1.

Using the MACs produced from the same simulated data using the proposed unmixing model the predicted mean abundance of each material in each satellite can also be computed using either the ANN or library matching unmixing models, and is shown in Figs. 22 and 23, respectively.

These results show particularly good agreement between the ground truth and proposed ANN model as to which materials are most abundant in each object. While both methods unmix the same general trend the library unmixing methods generally tend to over-estimate the maximum abundance observed for each material. However, it is rare for the unmixing models to predict 0% abundance of any material, instead often finding a relatively low abundance of the materials that are, in reality, not there. For example, most objects in the dataset contain GaAs due to the solar panels present on these satellites. Both models predict this well including correctly identifying relative abundances, such that Starlink satellites have around twice the amount of this material as Iridium-NEXT satellites. However, the two objects without GaAs have a prediction of 1.6% GaAs abundance with the ANN model, and higher in the library unmixing. This is higher than the true abundance of 0% but far lower than any object that really did have this material. Similar trends can be observed throughout this experiment.

For classification of satellites the critical factor will be how well our material probability estimates can distinguish false from true abundances of material, particularly for materials such as copper, where the true abundance to be detected is typically very low. As the ANN model achieves the most consistent performance, this model is used to estimate the probability of detection of each material. The probability of detection computed by the proposed model is shown in Fig. 24. Note that a probability score indicating how likely each material is to be present is computed for each instance of each object, with the mean score across the 2000 instances reported here. Thus 25% does not mean that the material is predicted 25% of the time but that on average that object has a 25% probability score of that material being present, which if consistent may result in the model never predicting that the material is truly there.

The material probabilities correlate well with prior expectations. Upper stages which are known not to contain solar panels, show a low probability of GaAs, as does the CubeSat that we do not contain panels. All satellites known to contain solar panels on the other hand show high probabilities of GaAs. So while there was some false positive abundance the probabilities clearly distinguish those objects that have this material from those that do not. Materials such as copper, that always exist in lower abundances in our dataset have greater uncertainty, with more




Figure 21. Ground truth mean abundance across 2000 samples of each satellite of a material being present on that satellite.

Unmixed Abundance

starlink	0.5658	0.2366	0.02134	0.05147	0.006672	0.01483	0.05622	0.03234	0.0147
iridiumNEXT	0.2466	0.2737	0.1045	0.05645	0.01379	0.01255	0.04354	0.0545	0.1944
HS376	0.6033	0.08568	0.09032	0.06008	0.005926	0.01683	0.0369	0.07778	0.02318
Upper Stage	0.01678	0.1393	0.5294	0.04499	0.01704	0.05019	0.03507	0.05607	0.1111
CubeSat(no panel)	0.0163	0.5035	0.0287	0.04195	0.3526	0.006248	0.005361	0.02163	0.02369
CubeSat(with panel)	0.2718	0.3352	0.02467	0.05638	0.2391	0.007772	0.008768	0.02418	0.03204
galileo	0.3994	0.03747	0.01884	0.04028	0.008942	0.007975	0.05012	0.01917	0.4177
calipso	0.616	0.0279	0.02317	0.09499	0.006622	0.006661	0.008309	0.01771	0.1985
glonassk	0.363	0.03343	0.1245	0.0638	0.01257	0.00939	0.02277	0.0327	0.3379
GPSIIF	0.4553	0.165	0.03626	0.0449	0.008395	0.01647	0.149	0.02781	0.09686
dubaiSat	0.09151	0.3643	0.02252	0.02977	0.01205	0.00552	0.01465	0.03397	0.4257
BoxWing	0.2951	0.0422	0.02603	0.04114	0.009918	0.004095	0.007879	0.02526	0.5484
dragon2	0.1853	0.04519	0.5874	0.0724	0.01521	0.01149	0.0248	0.04302	0.01513
	GaAs	Aluminium	White Paint	Black Paint	Green Paint	Red Paint	Copper	Titanium	Gold

Figure 22. Mean abundance of materials across 2000 samples of each satellite of a material being present on that satellite according to the ANN model.

Unmixed Abundance

starlink	0.7417	0.3413	0.1101	0.1156	0.03085	0.01043	0.1098	0	0.0164
iridiumNEXT	0.3904	0.4074	0.5667	0.04814	0.02584	0.005485	0.05588	0.02213	0.363
HS376	0.6757	0.4632	0.3946	0.519	0.02098	0.01172	0.1536	0.09554	0.08534
Upper Stage	0.0922	0.4801	0.9379	0.147	0.03435	0.1777	0.02268	0.5802	0.4203
CubeSat(no panel)	0.03441	0.6412	0.003214	0	0.3235	0.0001198	0.02055	0	0.0002317
CubeSat(with panel)	0.7371	0.6196	0.2051	0.2793	0.3137	0.00309	0.03226	0	0.006074
galileo	0.9055	0.2135	0.09855	0.0711	0.021	0.01545	0.3903	0	0.9651
calipso	0.941	0.15	0.06101	0.2213	0.02701	0.01764	0.06051	0.008097	0.3496
glonassk	0.9154	0.2211	0.7364	0.3956	0.03883	0.01716	0.09377	0.1284	0.6436
GPSIIF	0.9209	0.7413	0.4683	0.6034	0.03097	0.02607	0.6633	0.001321	0.292
dubaiSat	0.2254	0.7597	0.09512	0	0.0155	0.00783	0.01036	0.05229	0.6541
BoxWing	0.7638	0.1647	0.02439	0.0166	0.02094	0.003835	0.02951	0	0.9322
dragon2	0.6025	0.03277	0.9069	0.2718	0.01979	0.005439	0.04299	0.0005518	0.00135
	GaAs	Aluminium	White Paint	Black Paint	Green Paint	Red Paint	Copper	Titanium	Gold

Figure 23. Mean abundance of materials across 2000 samples of each satellite of a material being present on that satellite according to the library unmixing model.

	GaAs	Aluminium	White Paint	Black Paint	Green Paint	Red Paint	Copper	Titanium	Gold
starlink	0.9413	0.7613	0.2573	0.5954	0.1491	0.2455	0.5519	0.3273	0.1835
iridiumNEXT	0.9696	1	0.9389	0.8706	0.2597	0.2608	0.7374	0.5536	0.9393
HS376	1	1	0.9999	0.9867	0.1356	0.3295	0.9927	0.9121	0.6339
Upper Stage	0.2649	0.8854	0.9993	0.9327	0.2275	0.823	0.5977	0.8982	0.8189
CubeSat(no panel)	0.2681	1	0.2776	0.4212	0.9677	0.1797	0.1121	0.3281	0.2283
CubeSat(with panel)	0.9829	0.9955	0.3024	0.8891	0.9767	0.2339	0.1896	0.3881	0.4928
galileo	0.994	0.9268	0.3894	0.7379	0.2092	0.2473	0.8974	0.3708	0.9997
calipso	1	0.5393	0.4129	0.9671	0.1265	0.1156	0.1025	0.3097	0.9993
glonassk	0.9923	0.7286	0.9869	0.9569	0.3496	0.2914	0.74	0.5781	0.9981
GPSIIF	0.9995	0.9879	0.8874	0.9224	0.2454	0.3617	0.9881	0.5161	0.987
dubaiSat	0.9736	1	0.2783	0.4299	0.2265	0.1208	0.1858	0.4166	0.999
BoxWing	0.9958	0.6657	0.3893	0.5755	0.2173	0.1389	0.1223	0.3845	1
dragon2	0.9958	0.61	1	0.9206	0.2092	0.1274	0.3267	0.5225	0.1986

Figure 24. Mean probability across 2000 samples of each satellite of a material being present on that satellite according to the ANN model.

objects showing a mean probability of prediction closer to 50%. However in the majority of cases the probability scores correlate well with the objects that do have copper, with the exception of upper stages where false positive copper detections are relatively likely. In other cases the probabilities demonstrate uncertainty is likely to occur. Similarly, in the test dataset used here titanium exists only on HS376 and the Upper Stage. It is evident here that the probability of detection is significantly higher in these two examples, but there is also a reasonable probability of detection in a few other satellites, suggesting there are features there that might mimic titanium.

While it is useful to study the probability of materials present it is also useful to retain information of the abundance of each material, as shown in Fig. 22. When comparing this information it can be observed that while a similarly high probability of GaAs is found in Starlink and IridiumNEXT the Starlink satellite contains approximately twice as much of this material. Similar relationships can be assessed, with more or less accuracy, using each unmixing technique. For example, the ANN model will predict relatively accurately that Starlink has around five times as much of that GaAs material as the DubaiSat, while the library matching model underestimates this, at around three times as much. Even with a bit of variation in performance this still provides some useful information to understand the relative abundances of materials between objects to facilitate further analysis and classification. Abundance is estimated at every point in time though and in this instance it is the maximum abundance that is shown. It may be useful to study the mean abundance to assess the overall composition of the satellite, or to attempt to study particular faces of a satellite. For example, the time where GaAs abundance is maximised should indicate the solar panels facing towards the sensor. It may then be the case that it is known what materials should be visible from that perspective to help the classification.

Probabilistic object classification

The final stage in this work is to classify satellites into classes that describe their general composition or intended purpose. We use only the probability of each material being present on the object to minimise the dependency of the model on the satellites movement or the visibility of certain faces of the satellite. The previous section suggests that the ANN method produced more accurate results, thus the object classification method uses the probability of detection from only the ANN method as inputs. While the abundance of material present 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 3. Note that the Boxwing satellite is a generic shape so has no obvious class.

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

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 3. Full list of satellite models used and corresponding classes.

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 the materials that are expected to be present on each satellite, as shown in Fig. 25.

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 4 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 4. It can be observed from Fig. 24, 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

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

Figure 25. Materials that are expected to be present on each object in the simulated dataset.

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 4. List of classes and range of material probabilities used in synthetic training dataset.

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 paper 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 4. 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 Fig. 26. Note that the prediction rate for each object and class are rounded to 2 decimal places, as such the the sum of each row may equal slightly higher or lower than 1 in this and subsequent tables.

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. As significantly, the classification errors are typically logical and explainable. 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

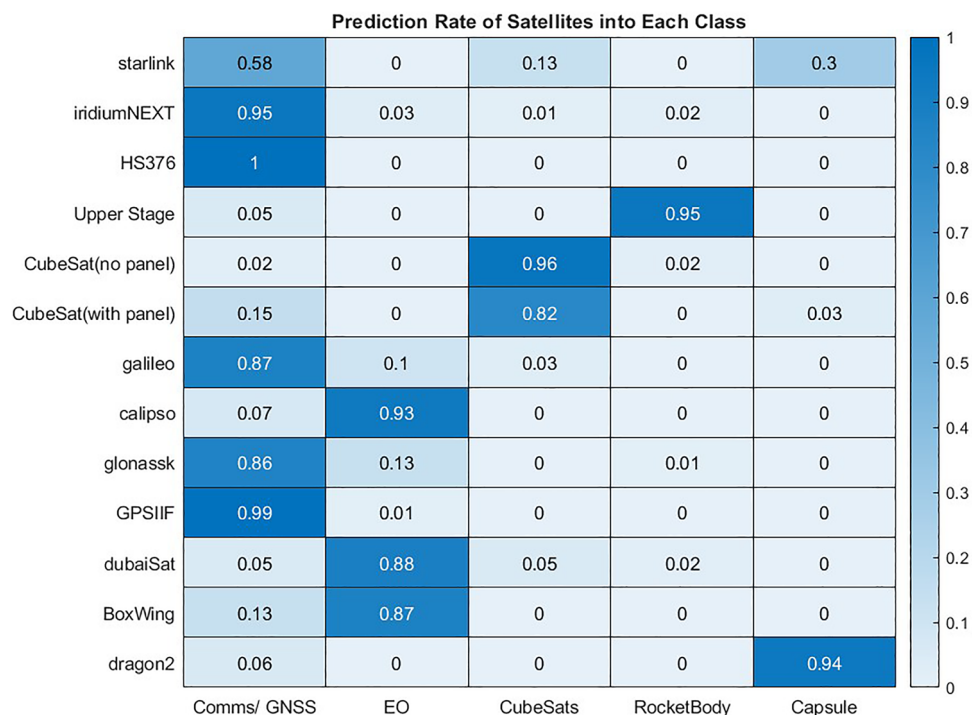


Figure 26. Prediction rate from the knn model when classifying objects based on probability of materials present on the object.

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, 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 than any other object known to contain it, as shown in Fig. 24. 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 result 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 Fig. 27.

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 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 used 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 in the 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




Figure 27. Mean probability of predicting object as belonging to each class in the KNN model.

classification. Any classification that does not match these assumptions are instead labelled as UFO. The revised classification results using this method are shown in Fig. 28. 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 Glonassk 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 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 it is more important to ensure those classified are correct, at the expense of failing to classify some at all. The probabilistic 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 classified with lower probabilities can be re-examined on multiple occasions to increase the information available and improve the confidence of the classification result.

Conclusion

In this paper we presented a novel data processing pipeline designed to determine the materials present on unknown space objects, and classify accordingly. This was comprised of a model to decompose a given hyper-spectral signal into the relative contribution from each present material, which was shown to be possible using an artificial neural network, or alternatively, more traditional techniques such as constrained least squares. The behaviour of these approaches was also investigated when applied in non-ideal scenarios; when there is discrepancy between the training set and the real materials (for example, with weathered satellites), and the case where the object contains a material that is not present in the library.

It was demonstrated that the material abundance curves can be used to estimate the probability of each material being present. A second model was then developed which estimates the probability of a space object belonging to predefined classes. Based on the output of this model satellites are labelled either as belonging to a known class or being unidentified and labelled as a UFO.

Generalisation to unseen object was prioritised in the design of these models, with both the decomposition and classification models being constructed in an inherently geometry-agnostic manner. Both models were trained using synthetic data that mimics the received spectra and output of the decomposition model respectively. Thus no space object data, real or simulated, was used to train the model. A validation campaign on 26,000 simulations split across 13 types of space objects, including simulation of atmosphere and randomised attitude, demonstrated good performance in predicting materials in classes.

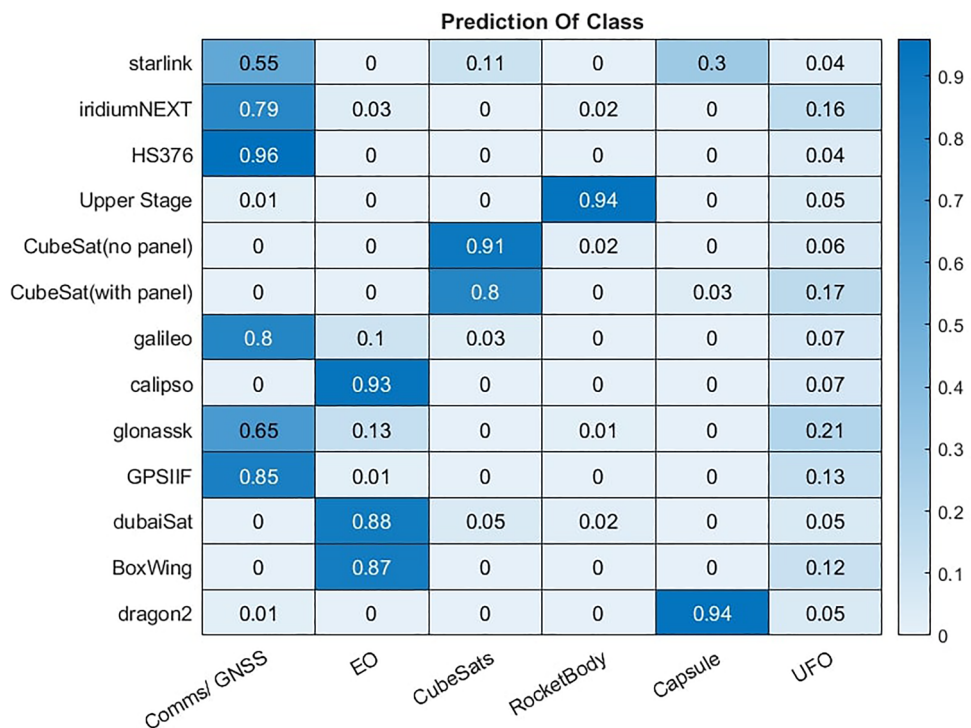


Figure 28. Prediction rate from the knn model when using model confidence and prior assumptions to flag uncertain predictions as UFO.

In cases where the object was classified incorrectly, we showed that the proposed pipeline was valuable, as by first analysing the materials present the reason for each classification decision could be easily understood and explained. It was found that typically errors are due to low abundances, that are harder to detect, or because the material exists on a small component of the satellite that is not visible at all during observation. In order to deal with the uncertainty in the classification process, we proposed a technique to derive a probability score from the machine learning model. The probability score captures the uncertainty in the prediction of the classes. By using two models that feed into each other, uncertainty in classification could be traced back to uncertainty in the presence of a certain material. This further enhances the explainability of this model.

Note that the definition of the classes in this paper serves only the scope to show how to group objects from their material composition. Thus it does not imply the specific function of any of the objects in each of the classes nor does it define what a EO or a comms object is. Alternative definitions are possible without changing the overall classification approach.

A complete classification system would, however, require an extended library, with proper consideration for material degradation, and a more in depth analysis of the resulting computational complexity. On the other hand, we showed that, for materials that appear in large abundance, the classification system working with a small library is quite robust and can provide a reliable association of objects to classes. When uncertainty is high and a deeper analysis is required the UFO class is effective at capturing undecided objects.

Furthermore, although spectra used in the simulations were measured in lab environment and we carefully modelled sensor noise, ageing and atmospheric effects, the ultimate proof of the effectiveness of the proposed methodology would require the acquisition of spectra of actual space objects with known composition. A partial validation with a known satellite component placed at a given distance from a telescope on Earth confirmed that if the materials are present in the library the component could be correctly identified and its material classified.

Data availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Received: 26 May 2023; Accepted: 8 January 2024

Published online: 18 January 2024

References

- Jorgensen, K. *et al.* Physical properties of orbital debris from spectroscopic observations. *Adv. Sp. Res. (Space Debris)* **34**, 1021–1025. <https://doi.org/10.1016/j.asr.2003.02.031> (2004).
- Abercromby, K., Guyote, M., Okada, J. & *et al.* Applying space weathering models to common spacecraft materials to predict spectral signatures. In *Proceedings of AMOS Technical Conference, Maui, Hawaii, USA* (2005).
- Reyes, J. A., Cowardin, H. M. & Cone, D. M. Characterization of spacecraft materials using reflectance spectroscopy. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii, September 11–14* (2018).
- Cardona, T., Seitzer, P., Rossi, A., Piergentili, F. & Santoni, F. Bvri photometric observations and light-curve analysis of geo objects. *Adv. Sp. Res.* **58**, 514–527. <https://doi.org/10.1016/j.asr.2016.05.025> (2016).
- Zhao, X.-F., Zhang, H.-Y., Yu, Y. & Mao, Y.-D. Multicolor photometry of geosynchronous satellites and application on feature recognition. *Adv. Sp. Res.* **58**, 2269–2279 (2016).
- Vananti, A., Schildknecht, T. & Krag, H. Reflectance spectroscopy characterization of space debris. *Adv. Sp. Res.* **59**, 2488–2500. <https://doi.org/10.1016/j.asr.2017.02.033> (2017).
- Willison, A. & Bédard, D. A novel approach to modeling spacecraft spectral reflectance. *Adv. Sp. Res.* **58**, 1318–1330. <https://doi.org/10.1016/j.asr.2016.06.013> (2016).
- Lind, L., Laamanen, H. & Pölonen, I. Hyperspectral imaging of asteroids using an FPI-based sensor. In *Sensors, Systems, and Next-Generation Satellites XXV* (Babu, S. R., Hélière, A. & Kimura, T. eds.). Vol. 11858. 65–78. <https://doi.org/10.1117/12.2599514> (International Society for Optics and Photonics, SPIE, 2021).
- Esposito, M. Spectral imagers for relative navigation, on-orbit servicing and debris removal. In *Clean Space Industrial Days* (ESA, 2018).
- Poojary, N., D'Souza, H., Puttaswamy, M. R. & Kumar, G. H. Automatic target detection in hyperspectral image processing: A review of algorithms. In *2015 12th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, 1991–1996 <https://doi.org/10.1109/FSKD.2015.7382255> (2015).
- Wetterer, C. J. & Jah, M. Attitude estimation from light curves. *AIAA J. Guidance Control Dyn.* **32**(5), 22 (2009).
- Dianetti, A. & Crassidis, J. Light curve analysis using wavelets. In *2018 AIAA Guidance, Navigation, and Control Conference, Kissimmee, Florida*, 8–12 January 2018 (2018).
- Matsushita, Y., Arakawa, R., Yoshimura, Y. & Hanada, T. Light curve analysis and attitude estimation of space objects focusing on glint. In *1st International Orbital Debris Conference (IOC), Sugar Land, Texas, USA* (2019).
- Chote, P., Blake, J. & Pollacco, D. Precision optical light curves of leo and geo objects. In *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference* (2019).
- Santoni, F., Cordelli, E. & Piergentili, F. Determination of disposed-upper-stage attitude motion by ground-based optical observations. *J. Spacecr. Rockets* **50**, 701–708. <https://doi.org/10.2514/1.A32372> (2013).
- Yanagisawa, T. & Kurosaki, H. Shape and motion estimate of leo debris using light curves. *Adv. Sp. Res.* **50**, 136–145. <https://doi.org/10.1016/j.asr.2012.03.021> (2012).
- Kerr, E. *et al.* Light curves for geo object characterisation. In *8th European Conference on Space Debris* (2021).
- McNally, K., Ramirez, D., Anton, A., Smith, D. & Dick, J. Artificial intelligence for space resident objects characterisation with light curves. In *8th European Conference on Space Debris* (2021).
- Vasile, M. *et al.* Intelligent characterisation of space objects with hyperspectral imaging. *Acta Astron.* **203**, 510–534. <https://doi.org/10.1016/j.actaastro.2022.11.039> (2023).
- Vasile, M. *et al.* *Hyperspectral Classification of Space Objects* (2022).
- Clark, R. N. *et al.* *USGS digital spectral library splib06a* (Technical Report, US Geological Survey, 2007).
- Winter, M. E. N-findr: An algorithm for fast autonomous spectral end-member determination in hyperspectral data. In *Imaging Spectrometry V*. Vol. 3753. 266–275 (SPIE, 1999).
- Harsanyi, J. C. & Chang, C.-I. Hyperspectral image classification and dimensionality reduction: An orthogonal subspace projection approach. *IEEE Trans. Geosci. Remote Sens.* **32**, 779–785 (1994).

24. Nascimento, J. M. & Dias, J. M. Vertex component analysis: A fast algorithm to unmix hyperspectral data. *IEEE Trans. Geosci. Remote Sens.* **43**, 898–910 (2005).
25. Heylen, R., Burazerovic, D. & Scheunders, P. Fully constrained least squares spectral unmixing by simplex projection. *IEEE Trans. Geosci. Remote Sens.* **49**, 4112–4122 (2011).
26. Zhang, Z. & Aeron, S. Exact tensor completion using t-svd. *IEEE Trans. Signal Process.* **65**, 1511–1526 (2016).
27. Nie, B., Yang, L., Zhao, F., Zhou, J. & Jing, J. Space object material identification method of hyperspectral imaging based on Tucker decomposition. *Adv. Sp. Res.* **67**, 2031–2043 (2021).
28. Aggarwal, H. K. & Majumdar, A. Hyperspectral unmixing in the presence of mixed noise using joint-sparsity and total variation. *IEEE J. Sel. Top. Appl. Earth Observ. Remote Sens.* **9**, 4257–4266 (2016).
29. Duzellier, S. *et al.* Space debris generation in geo: Space materials testing and evaluation. *Acta Astron.* **192**, 258–275 (2022).
30. Pearce, E. C., Weiner, B. & Krantz, H. Examining the effects of on-orbit aging of SL-12 rocket bodies using visible band spectra with the MMT telescope. *J. Sp. Saf. Eng.* **7**, 376–380 (2020).

Acknowledgements

This project is supported by the ESA OSIP grant HyperClass. The authors would like to thank Emiliano Cordelli and Jan Siminski of the ESA debris office for their advice and guidance.

Author contributions

M.V. conceived the idea and methodology, directed the implementation of the pipeline and the execution of the experiments, analysed the results and reviewed the paper, L.W. implemented the simulation model and the machine learning material unmixing, A.C. developed the probabilistic material and object classification, S.M. supported the implementation of the simulation model, V.S. extracted the spectra from the laboratory experiments, P.M. and S.M. provided key input to the development of the simulations and the analysis of the results. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.V.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024