

Convolutional Neural Networks for Inference of Space Object Attitude Status

GREGORY BADURA,¹ CHRISTOPHER R. VALENTA,¹ BRIAN GUNTER,²

¹*Georgia Tech Research Institute, Atlanta GA*

²*Georgia Institute of Technology, School of Aerospace Engineering, Atlanta, GA.*

ABSTRACT

In this study, convolutional neural networks (CNNs) are utilized for the task of inferring the attitude status in terms of rotation rate of resident space objects (RSOs) using simulated light curve measurements. Research into the performance of CNNs on synthetic light curve data-sets has shown significant promise that has not yet translated into success when working with empirically collected light curves. This limitation appears to be due to a number of factors including: mixing of bidirectional reflectance distribution function (BRDF) signatures, the effects of sensor noise, and blurring due to atmospheric turbulence. A synthetic data-set of approximately 7500 light curves was generated that takes into account realistic BRDF signatures and environmental parameters. The RSO used in this study was texture mapped with three unique material BRDF signatures: silicon solar panel, glossy paint, and aluminum. A two-step BRDF model inversion of the Beard-Maxwell model was performed using empirically collected data-sets of these materials in order to physically derive the BRDF model parameters. The CNN was trained on light curves resulting from the RSO performing four different maneuvers: tumbling, accelerating in rotational rate, stabilizing, and inactive (or stable in rotation rate). The CNN achieved an overall classification accuracy of 86.2% across the four maneuver classes. A confusion matrix analysis of the different classes of maneuvers suggested that our model performed best when classifying tumbling and accelerating RSOs (94% accuracy) and worst at classifying inactive RSOs (60% accuracy). This performance limitation when classifying inactive RSOs to (1) back-scatter signatures and specular glints within the synthetic light curves of inactive satellites being mistaken as attitude maneuvers, and (2) low signal-to-noise ratio due to factors such as atmospheric blurring. These results suggest that CNNs have strong potential for aiding in the problem of classifying satellite attitude status from light curves, but that machine learning research must focus on developing training sets and pre-processing techniques that account for these complications.

Keywords: machine learning – Convolutional Neural Networks (CNNs) – bidirectional reflectance distribution function (BRDF) – light curve

1. INTRODUCTION

Inferring a satellite’s operational status and attitude is a critical task in the field of space domain awareness (SDA) (Holzinger & Jah 2018; Spurbeck et al. 2018). Observing the visible photometric light curve of a resident space object (RSO) can enable the ability to track its movements and ensure it is maintaining attitude. With sufficient awareness of a nearby RSO’s attitude state, a RSO could potentially perform avoidance maneuvers to avoid collisions with disastrous consequences (DiBona et al. 2019). In order to enhance awareness of potential threats, researchers need to develop algorithms that can predict the operational status of RSO’s that neighbor each other and maintain a catalogue of any changes in attitude state. By monitoring a sufficiently large region, researchers could provide sufficient warning to maneuverable RSO’s to engage in predictive avoidance.

Current methods of predicting a RSO’s attitude state often rely on simultaneously inferring the shape and attitude of the RSO of interest (Früh et al. 2013; Wetterer & Jah 2009; Fulcoy et al. 2012). In these methods, the RSO is

Corresponding author: Gregory Badura
gregory.badura@gtri.gatech.edu

often assumed to be a Lambertian reflector or the RSO is assigned an arbitrary bi-directional reflectance distribution function (BRDF) in order to describe its photometric behavior.

There are two potential problems with these approaches. Researchers have shown that satellite materials have unique and spectrally varying BRDF behaviors that physically un-realistic BRDF models fail to capture (Willison & Bédard 2016; Reyes & Cone 2018). In addition, current methods of performing light curve inversion on attitude and shape simultaneously can struggle when the shape is non-convex, yet many RSO's do not meet this stringent shape criteria due to the presence of extruding solar panels and antennae (Bradley & Axelrad 2014).

An example of the variability in RSO reflectance magnitude as a function of pose can be seen from hyperspectral camera measurements obtained from the Georgia Tech Ranging and Nanosatellite Guidance Experiment (RANGE) satellites in Figure 1 (c) (Gunter et al. 2016). RANGE consists of two 1.5-unit sized cubesats deployed in a leader-follower formation as they orbit the earth in low-earth orbit (Figure 1 (a)). The two satellites are designed to communicate both with each other and a ground support team at Georgia Tech. The satellites' surfaces are comprised of a variety of materials with unique BRDF signatures including metals, paints, and solar panels as can be seen in the color-coded diagram in Figure 1 (b). Hyperspectral camera measurements of the satellite obtained under diffuse illumination conditions show that the magnitude of the reflectance changes significantly depending on the orientation of the satellite relative to the observer.

Current research in the SDA has shown that machine learning techniques could be invaluable to inferring the attitude status of RSOs from light curves. For example, Dao et al have shown that machine learning algorithms can achieve up to 87% classification accuracy when labeling light curves of geo-stationary objects as either a tumbling rocket body or tumbling satellite (Dao et al. 2019). In another study, Furfaro showed that machine learning techniques can achieve 92% accuracy in solving the problem of RSO light curve inversion for shape retrieval (Furfaro et al. 2019). These early successes suggest that machine learning techniques could potentially help solve the problem of inferring an RSOs attitude status from its light curve signature.

One particular machine learning algorithm that could be of interest for the problem of inferring a satellite's attitude status from light curves is the convolutional neural network (CNN) which has already shown significant promise in fields ranging from image classification (Marana et al. 1997; Parisi et al. 1998) to time-series classification of medical data-sets (Haselsteiner & Pfurtscheller 2000; Wang et al. 2017). Recently, CNNs have begun to experience popularity in the SDA field in accomplishing tasks such as space object classification from light curve datasets (Furfaro et al. 2018; Linares & Furfaro 2016), shape retrieval from light curve data (Furfaro et al. 2019; Fan et al. 2019), and orbital path prediction (McQuaid et al. 2018; Peng & Bai 2019).

CNNs have shown excellent performance when being trained and tested on synthetic light curve datasets of RSOs textured with a single BRDF signature. Linares & Furfaro (2016) trained a CNN to classify light curves from a synthetic dataset into one of four categories: rocket bodies, controlled payloads, uncontrolled payloads, and debris. The RSOs had a single BRDF signature that was not published in the paper. Furthermore, these RSOs traveled in identical geo-stationary orbits, which simplified the task of dealing with observer-solar phase angle effects (Hall et al. 2007). The CNNs trained on this synthetic dataset achieved a classification accuracy of 99.6%, indicating the potential utility of CNNs in the tasks of RSO shape and attitude estimation.

Despite these signs of success, CNNs have struggled when trained and tested on light curves from empirical measurements of RSOs that are comprised of a mixture of BRDF signatures. In a follow up study, Furfaro et al. (2018) extended the CNN to be trained on synthetic light curves from the study performed by Linares & Furfaro (2016) and on empirical light curves of RSOs obtained by the Multichannel Monitoring Telescope (MMT). These empirical light curves obtained by the MMT were more clearly afflicted by noise and atmospheric effects. It also appeared that the rocket body signatures had occasional specular glints that did not occur in the synthetic light curve dataset. The CNN performed far worse when tested on the real light curve data-sets, achieving just 75% accuracy versus 99.6% on the synthetic data-set.

Clearly, there is a need for testing the performance of CNNs when classifying RSO attitude status on synthetic light curves with system attributes such as: physically-derived BRDF signatures, multiple materials coating the surface of the RSO, atmospheric modeling, phase-angle effects, and sensor noise characteristics. In this paper, we evaluate the performance of a CNN that was trained using synthetic light curves derived using these factors. This paper makes several novel contributions to the field of SDA, including: (1) the incorporation of realistic mixed BRDF signatures into modeling the light curve of an RSO, (2) the derivation of a light curve modeling approach that takes into account

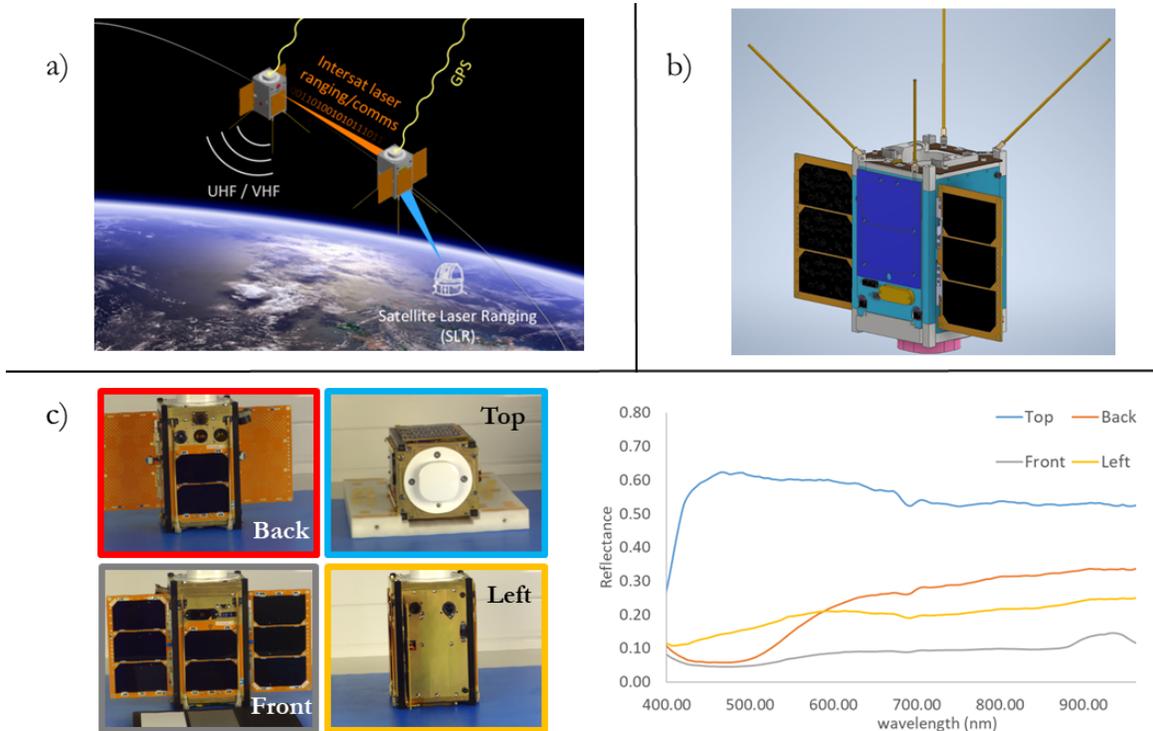


Figure 1. Characterization of the Georgia Tech launched RANGE-A cubesat: (a) a diagram of the "leader-follower" formation of the cubesat in orbit, (b) a CAD model of RANGE-A diagramming the various facets of the satellite, and (c) a laboratory hyperspectral characterization of the satellite obtained prior to launch.

the effects of atmospheric blurring and sensor sampling, and (3) the presentation of results suggesting strengths and weaknesses of CNNs in the task of ascertaining attitude status from light curves.

The format of this paper takes on the following format. In Section 2, we describe the theory behind CNNs and present the approaches that were used to train the architecture used in this paper. In Section 3, we present the derivation of a light-curve simulation that takes into account effects such as motion-induced smearing of the signal and turbulence blurring of the RSO signature. In Section 4, we describe the attitude maneuvers that serve as the basis for the classes that the CNN was trained on: tumbling, accelerating, stabilizing, and inactive. In Section 5, we present the datasets, inversion procedure and metrics that were used to derive physically realistic BRDF signatures that can describe the reflectance characteristics of the RANGE satellite. Finally, in Section 6, we present the resultant light curves and the performance of our CNN in classifying these light curves according to the true attitude maneuver.

2. CONVOLUTIONAL NEURAL NETWORKS

The popularity and strength of CNNs lies in the automation of feature extraction that is performed in the process of training the CNN to discriminate among different classes. CNNs have been covered in the literature extensively both in the field of SDA (Linares & Furfaro 2016; Furfaro et al. 2018) as well as in the field of computer vision (Krizhevsky et al. 2012; Simonyan & Zisserman 2014). For this reason, only the basics of the layers comprising the CNN architecture and the process of training a CNN will be described here.

2.1. Theory of CNN

The input light curve data set to the CNN consists of light curve measurements denoted by time-series vectors $\mathbf{x} \in \mathbb{R}^{T_s \times 1}$, where T_s is the time-series length of the light curve vector. The output class vector is an integer representing one of N_c classes of maneuvers denoted by the vector $\mathbf{y} \in \mathbb{R}^{N_c \times 1}$. A CNN can be thought of as composing two phases: a training phase and a prediction phase. During the training phase, the deep neural network learns to map from light-curve vector \mathbf{x} to class vector \mathbf{y} using a subset of the total number of labeled light curve vectors. During the prediction phase, the deep neural network makes predictions on unseen light curve vector using the feature vectors that it has learned in the process of training (Linares & Furfaro 2016).

During the training phase, CNNs learn a parameterized mapping from signal time-series to relevant features via a process of gradient descent on batches of a training dataset. The end-result of the training procedure is a parameterization of the weights of convolution kernels, which are ultimately convolved with the time-series datasets in order to produce feature maps (McQuaid et al. 2018). The total number of feature maps produced can vary based on the total number of convolution layers used in the network architecture, in order to achieve a balance between over-fitting and under-fitting the CNN on the training data. The weights used in the convolution kernels are trained via gradient descent to minimize a cost function between the true classes of the training data and the predicted classes of the training data by the CNN. Convolution layers can be stacked to increase the feature-learning capacity of the CNN. This increases the computational requirements, but operations such as pooling can be used to reduce the dimensionality in between convolution layers (Zhao et al. 2017).

2.2. Architecture of CNN

Each convolution layer, $\mathbf{h}^{cov}(\mathbf{x})$, used in this study has the following form (Linares & Furfaro 2016):

$$\mathbf{h}^{cov}(\mathbf{x}) = \mathbf{f}(\mathbf{x} * W^{cov} + b^{cov}) \quad (1)$$

where $*$ denotes the convolution operator, W^{cov} denotes the convolution kernel, b^{cov} is the bias of the convolution filter, and \mathbf{f} is the activation function for each layer that adds non-linearity to the feature vector. The activation function chosen for this study is the *ReLU* function which is denoted by $\mathbf{f}(\mathbf{i}) = \max(0, \mathbf{i})$, such that the function maps input vectors less than zero to negative input values in order to add non-linear features to the model. A CNN typically applies a series of these convolution kernels W^{cov} in layers, where each layer has a different width kernel that extracts features at varying scales (Zhao et al. 2017).

Several enhancements are made to the convolution layers in order to avoid the problem of overfitting to our training datasets. In this study, an activity regularization term is included in the training of the convolution layers that is selected via hyper-parameter tuning to improve generalizability of the model and reduce training error (Krizhevsky et al. 2012; Li et al. 2019). Batch-normalization (BN) is applied to the output of each convolution layer's activation function based on the demonstrated ability of this procedure to improve training speed and produce more stable gradient behavior (Santurkar et al. 2018). After BN, a max-pooling operation is performed in which the a feature map is divided into equal-length segments of size 4×1 , and then every segment is represented by its maximum value. The advantage of this pooling operation comes in the form of down-sampling the convolution layer output bands, reducing variability in the hidden layer activation function outputs (Swietojski et al. 2014).

The final layers of the CNN used in this study is the fully-connected layer. Fully connected layers are designed to connect every neuron in one layer to every neuron in another layer. The fully connected layer, $\mathbf{h}^{fc}(\mathbf{x})$, can be represented mathematically by (Linares & Furfaro 2016):

$$\mathbf{h}^{fc}(\mathbf{x}) = \mathbf{f}(W^{fc}\mathbf{x} + b^{fc}) \quad (2)$$

where the convolution operator in Equation 1 has been replaced by a multiplication operator with the term W^{fc} . The term b^{fc} represents the bias weight of the fully-connected layer.

The activation function of the final fully-connected layer is defined by a *softmax* function and is used to provide probability values that the light curve belongs to the N_c classes represented by values in the ranges $(0, 1)$ that sum to 1. This function is normally employed as the activation function at the output layer in order to represent a probability distribution over a discrete variable with N_c possible values (Goodfellow et al. 2016).

A diagram of these layers in the CNN architecture used in this study is shown in Figure 2. It was inspired by the architecture of the LeNet CNN (LeCun et al. 1998) and tuned using hyperparameters with the goal of achieving optimal performance on our datasets. The architecture will be described further in Section 6.3.

2.3. Training of CNN

The parameters which defined the convolution layers and the fully connected layers are in this case encased in a vector Θ . The goal of training the CNN is to choose a set of parameters Θ such that a loss function defining the error between true labels and predicted labels of the training data set is minimized. The loss function defined in this paper is the cross-entropy loss between true labels and predicted labels (Linares & Furfaro 2016):

$$\mathbf{L}(\Theta) = \frac{1}{N_l} \sum_{i=0}^{N_l} [\mathbf{y}_i \log \hat{\mathbf{y}}_i + (1 - \mathbf{y}_i) \log(1 - \hat{\mathbf{y}}_i)] \quad (3)$$

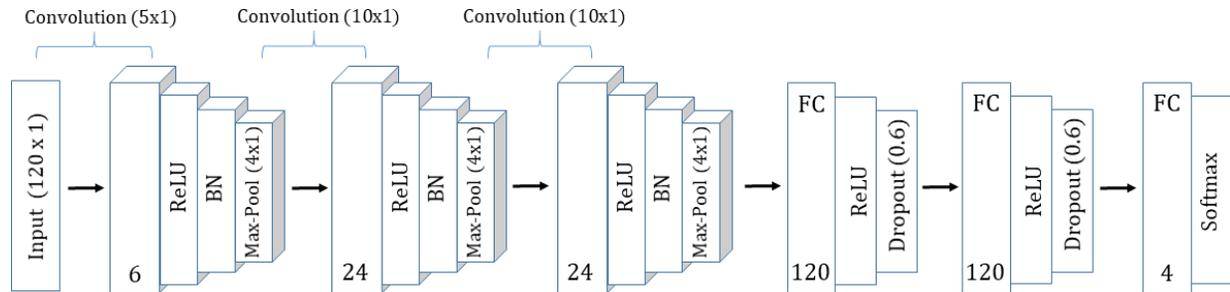


Figure 2. The CNN architecture resulting from hyper-tuning the parameters in this study. It employs the various layers described in Section 2.2. The numbers denote the number of filters per convolutional layer, and the size of the fully connected layer.

where N_i is the total number of light curves in the dataset, $\hat{\mathbf{y}}_i$ is the class label predicted by the model for the i^{th} light curve at a given epoch iteration, and \mathbf{y}_i is the true class label.

The Adam optimizer was used in order to find the global minimum solution to Equation 3. Adam differs from standard gradient descent in that the weight updates and the learning rate is allowed to adaptively change during training. In addition, a learning rate is maintained for each network weight and separately adapted as the training progresses. The algorithm has been shown to have favorable performance relative to other stochastic optimization routines (Kingma & Ba 2014).

In order to avoid the overfitting the parameters of Θ to our training data-set, early-stopping was implemented by monitoring the value of the loss function in Equation 3 on a testing-dataset that was not used to train the model. This procedure is based on previous studies showing that training can be stopped at the point of smallest error with respect to a validation data set (Caruana et al. 2001). We assigned an empirically selected patience parameter of 45 epochs in order to allow for the potential that the gradient descent routine was trapped in a local minimum.

3. LIGHT CURVE SIMULATION

A suite of tools for simulating light curves from RSOs in low-earth orbit was developed as a part of a research effort for studying the operational status of LEO satellites. The model was inspired by previous research on studying the detection-limits of telescopes located in urban environments with sky backgrounds that are degraded by light pollution (Badura et al. 2019). The model is capable of simulating time varying light curves for a modeled optical telescope at a defined geographic location of interest. It ingests various parameters in order to simulate a time-varying light curve, including: orbital characteristics as defined by a Two-Line Element (TLE) set, the rotational state of the RSO body, the atmospheric turbulence strength, the telescope optical characteristics, the shape of the RSO body, the directional reflectance characteristics of the RSO's surface materials, and the sky background brightness. For the sake of focusing on the machine learning results of our study, the model is fully derived in Appendix A.

4. SATELLITE ROTATIONAL AND ORBITAL MECHANICS SIMULATIONS

In this study, the Cartesian coordinate system of the CAD model for the RANGE-A satellite is used as the Cartesian coordinate system of the body frame of the satellite in orbit. The orientations and positions of the sun, the observer, and the outward facet normals are updated in time relative to this body frame according to the Two-Line Element (TLE) set derived orbital mechanics of the satellite and the rotational mechanics of an attitude maneuver.

4.1. Satellite Body Frame Overview

The directional unit vector $\hat{n}_j^{B,0}$ denotes the unit vector specifying the default outward normal of the j^{th} facet within the CAD model's coordinate system. The facets rotate relative to this Cartesian coordinate system according to defined maneuvers such that the orientation of the j^{th} facet relative to the coordinate system changes over time. The unit vector denoting the direction of the j^{th} facet within the body frame at a given time t is denoted by the parameter $n_j^B(t)$.

In this study, the orientation of the j^{th} facet at the initial observation of the light track, $n_j^B(t=0)$, was calculated by performing a random rotation of $\hat{n}_j^{B,0}$ within the body frame to simulate a random initial orientation of the facet relative to the observer and sun at the start of the observed light curve. The orientation for times $t > 0$ was determined

by using a randomly selected rotation axis, and a time series of angular velocities and accelerations meant to simulate an attitude maneuver. The attitude maneuver was represented by a time varying Euler-Rodrigues rotation matrix that represents the rotation relative to the initial position of the satellite at $t = 0$ according to one of the following maneuver categories: (1) stabilizing, (2) freely tumbling, (3) accelerating, and (4) inactive.

The categories of stabilization maneuvers are defined by the following equations for the angular velocity of the satellite about a randomly selected rotation axis:

1. For the case of a stabilizing satellite, the attitude sequence took on on the form of either a satellite that is achieving constant stabilization or one that is oscillating throughout the period of stabilization due to competing forces.

- (a) The **oscillating stabilizing** satellite is represented by a decreasing angular velocity according to a dampened harmonic oscillator:

$$\omega_{d,s}(t) = \omega_0 \exp\left(\frac{\ln(2)t}{\gamma}\right) \cos\left(\left(\frac{2\pi}{T}\right) t\right), \quad (4)$$

where ω_0 represents the initial angular velocity of the satellite about the rotation axis, γ represents the half life dictating the rate at which $\omega_{d,a}$ decays over time, and T represents the period of oscillations for the satellite.

- (b) The **constant stabilizing** satellite is represented by the following equation:

$$\omega_{c,s}(t) = \omega_0 + \alpha_{a,0}t \text{ for } t < \frac{-\omega_0}{\alpha_{a,0}}, \quad (5)$$

where $\alpha_{a,0}$ represents the initial angular acceleration of the satellite about the rotation axis.

2. For the case of a **freely tumbling** satellite, the satellite is represented by a constant angular velocity $\omega_t(t) = \omega_0$ with angular acceleration $\alpha_t = 0, \forall t$.
3. For the case of an **accelerating** satellite, the attitude sequence takes on the same forms of the stabilizing satellite, except for the time series of Equations 4 and 5 are reversed in time. The notation for the angular velocity of a constantly accelerating and dampened accelerating satellite maneuver are denoted by $\omega_{c,a}(t)$ and $\omega_{d,a}(t)$, respectively.
4. For the **inactive** satellite, both the angular acceleration and angular velocity are zero for all times throughout the light curve sequence such that the positions of the facets of the satellite relative to the body frame do not change.

Some example time series of the rotation maneuvers over a period $T_s=120$ seconds are shown in Figure 3. The angular velocity and the angular acceleration (calculated by $\delta\omega_i/\delta t$) of the i^{th} attitude maneuver are demonstrated for randomly generated sequences used in this study.

There are a number of ways to specify an attitude maneuver of a satellite employing both Euler angles and quaternions (Shuster 1993). The method chosen to specify the attitude of the facets at a given time is the Rodrigues rotation formula. The formula is a method for rotating the normal vectors of the facets relative to the body frame about a randomly chosen axis of rotation, \hat{k}_{rot} , and an angle of rotation (Dai 2015). The rotation matrix specifying the time varying orientation of the j^{th} facet relative to the original orientation of the facet at $t = 0$ is given by:

$$\mathbf{R}_{att}(t) = \mathbf{I} + \sin \theta_i(t) \left[\hat{k}_{rot} \times \right] + (1 - \cos \theta_i(t)) \left[\hat{k}_{rot} \times \right]^2, \quad (6)$$

$$\left[\hat{k}_{rot} \times \right] = \begin{bmatrix} 0 & -k_z & k_y \\ k_z & 0 & -k_x \\ -k_y & k_x & 0 \end{bmatrix},$$

where \mathbf{I} is the identity matrix, $[\hat{k}_{rot} \times]$ is the cross-product matrix of the axis of rotation, and $\theta_i(t)$ is the rotation angle due to the attitude maneuver that is specified by the time sequences of angular velocity and acceleration. Note that this is a simple treatment of attitude that does not incorporate terms such as solar-radiation pressure or inertial effects on the rotating body (Linares & Furfaro 2016).

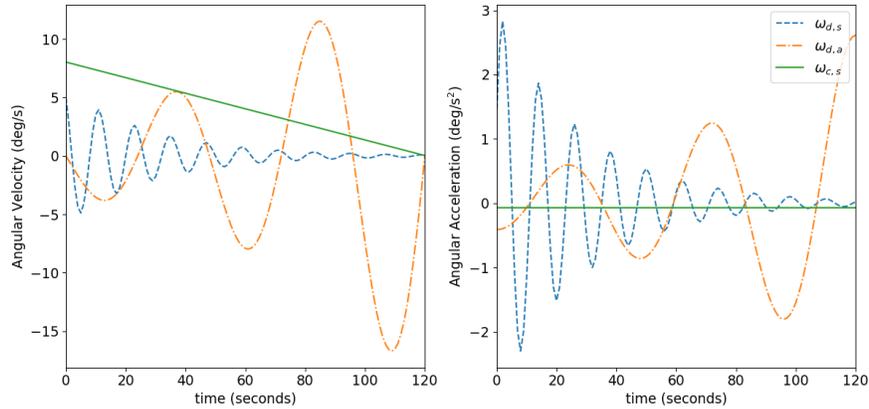


Figure 3. An example of 3 series of randomly generated rotation sequences with the angular velocity (left) and the angular acceleration $\delta\omega/\delta t$ (right): (blue) an oscillating stabilizing sequence with $T = 12$ seconds, $\omega_0 = 5$ deg/s, and $\gamma = 20$ seconds; (orange) an oscillating accelerating time sequence with $T = 48$ seconds, $\omega_0 = 24$ deg/s, and $\gamma = 45$ seconds; (green) a constantly stabilizing satellite with $\omega_0 = 10$ deg/s and $\alpha_{a,0} = 0.067$ deg/s².

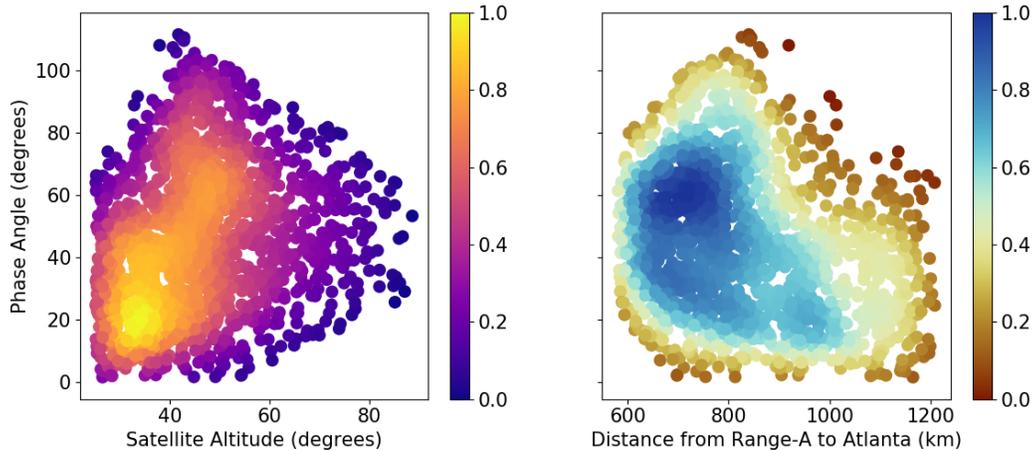


Figure 4. Normalized heat-maps showing the phase angle of the solar-observer geometry plotted as a function of satellite altitude above the horizon in Atlanta, GA (left) and distance from the satellite to the observer (right).

4.2. Orbital Mechanics Simulation Overview

Time series of position data were derived using TLE data for RANGE-A over a period of 5 years spanning from 2015 to 2020. At each moment in time, the instantaneous position of RANGE-A in the Geocentric Celestial Reference System (GCRS) was calculated. The GCRS position of RANGE-A was then converted into zenith and azimuth coordinates relative to the horizon of Atlanta, GA using the Astropy Python module (Price-Whelan et al. 2018). Tracks where RANGE-A had an altitude angle of greater than +30 degrees were deemed to be visible by a ground-based telescope. This stipulation was made based on the extreme blur when transmitting an optical wavefront horizontally through a turbulent atmosphere (Lawrence & Strohbehn 1970). The approximately 2000 tracks meeting this altitude criterion were designated as acceptable tracks and saved. The time series of GCRS positions of the sun for these tracks was then calculated using the Astropy python module (Price-Whelan et al. 2018). Using the law of cosines, the geometry of the observer and sun relative to the RANGE-A body frame was then calculated for these time-series tracks. The resultant derived parameters \hat{s}^B and \hat{o}^B denote the unit vectors specifying the direction from the satellite to the sun and observer, respectively, relative to the body frame. Figure 4 shows normalized heat-maps of the broad range of phase angles, satellite-to-observer distances, and satellite altitudes for the derived RSO orbital geometries.

5. BIDIRECTIONAL REFLECTANCE DISTRIBUTION FUNCTION MODELLING

The RANGE-A satellite CAD model is composed of upwards of 20 unique materials, with many of them color-coded in the diagram of Figure 1 (b). These include a variety of metals, paints, and composite materials. Given the lack of publicly available BRDF information on all of these materials, we chose to down-sample the number of different materials used in our RSO model to just three: alodined aluminum, silicon solar panels, and white paint. The sources of these datasets and the measurement procedures used to acquire them are discussed in Section 5.1.

The facets of the RANGE-A model were labeled with silicon solar panel BRDF and white paint BRDF for facets covered by these materials in Figure 1 (b). For all other facets, the aluminum BRDF data was used. The wavelength used in fitting the BRDF data to the theoretical model discussed in Section 5.2 was chosen to be 600 nm.

In the following three sections, we describe the sources of the BRDF data, the theoretical BRDF models to which the empirical BRDF data were fit, and the model inversion results that are used to model the material BRDF outside of the empirical view-illumination geometries.

5.1. Material Reflectance Data

The data used to represent the BRDF properties of the samples used in this study were acquired from various published sources of directional reflectance. These sources were chosen based on the perceived accuracy with which the measurements were obtained in addition to their ability to represent the characteristics of materials that are used on spacecrafts. A brief summary of the materials is shown in Table 1. The instruments used to acquire the sources and the number of viewer-illumination geometries extracted from the dataset (N_m) are also shown in the Table 1.

In order to represent the BRDF of silicon solar panels, data acquired by Hou et al. (2019) of space-based silicon solar panels matching the characteristics found in a study by Zhang et al. (2018) were used. The study sample was not solely of the solar cell but rather the entire solar panel including the cover glass, ethylene-vinyl acetate (EVA), silicon solar cells with electrode and back glass. The authors created a unique goniometric system that propagated light from a solar simulator to a test sample that was mounted on a rotating turntable. The incident light onto the sample was then reflected to an illuminometer situated on a semi-circular track. The illuminometer rotated about the track in order to simulate different solar-observer geometries.

The white paint sample used to represent the coating of the GPS receiver on RANGE-A was acquired from Cornell's Bidirectional Reflectance database (Cornell 2002). BRDF data was measured using a custom-built goniometer system to comply with ASTM Standard E1392-96. For each pair of incident and reflection directions, polarized measurements were acquired and combined into unpolarized reflectance data. The goniometer was equipped with a CCD camera system and one of 80 spectral filters that were used to obtain wavelength-dependent measurements (Marschner et al. 2000). The sample chosen for use in our study was a shiny paint that is applied to car surfaces.

The aluminum sample chosen for this study was from measurements of aluminum 6061 that would normally be used to assemble cubesats (Bédard et al. 2015). The spectrometer used to acquire the measurements was a Analytical Spectral Devices (ASD) FieldSpec 4 which covers the spectral range of 350 nm to 2500 nm. The light source was a 250 Watt halogen lamp placed at the focal point of a 15-cm parabolic mirror to produce a narrow cone of illumination onto the sample (Bedard et al. 2011). We note that there are a variety of factors including the anodization and polishing of an aluminum sample that can affect the specular and diffuse reflectance characteristics that were not mentioned in the original study and are not considered here.

Given the variety of approaches and measurement setups that were used to acquire the data shown in Table 1, it is clear that some standardization of measurement criteria should be adopted by the SDA community in the future. However, given the paucity of publicly accessible BRDF measurements of satellite materials and the goal of this study, we do not cover that topic in this paper.

5.2. Beard-Maxwell BRDF Model

The Beard-Maxwell BRDF model is based on a treatment of surfaces as being composed of randomly oriented microfacets leading to optical phenomena such as self-shadowing and geometric obscuration. A geometrical optics treatment is assumed in the model such that each microfacet behaves as a specular reflector obeying Snell's law of reflection. The model was originally developed to describe the reflection properties of rough, painted surfaces displaying Fresnel effects but was later extended to a wider variety of surfaces. The model assumes that the material surface is a three dimensional terrain of microfacets of varying orientation that coat a macro-facet (Maxwell et al. 1973). The model has shown significant success in being fit via model inversion to empirical measurements of a wide range of

Table 1. A summary of the BRDF database sources for the materials used to model the directional reflectance of RANGE-A and the information on the instruments used to measure the directional reflectance of each material.

| Material | Source | Summary of Sample Properties | Number of Orientations (N_m) | Spectral Measurement Device | Spectral Range |
|---------------------|-------------------------|--|----------------------------------|----------------------------------|------------------|
| Silicon Solar Panel | Hou et al. (2019) | The effective BRDF of the solar panel was measured including the cover glass, ethylene-vinyl acetate (EVA), silicon solar cells with electrode and back glass. | 160 | TES-1339 Illuminometer | 400 nm - 700 nm |
| Aluminum | Bédard et al. (2015) | The authors only stated that they studied 6061-T6 aluminum alloy, with no reference to the treatment of the surface in terms of anodization or polishing. | 300 | ASD FieldSpec 4 Spectrometer | 400 nm - 2500 nm |
| White Paint | Marschner et al. (2000) | The samples are representative of glossy paint that would be applied to cars. | 1500 | Photometrics PXL1300L CCD camera | 400 nm - 700 nm |

materials such as: metals, manufactured materials, soils, and water (Westlund & Meyer 2002). The Beard-Maxwell model used in this study can be written in units of inverse steradians as (Westlund & Meyer 2002; Montanaro 2007)

$$\rho_{BM}(\theta_i, \phi_i; \theta_r, \phi_r) = R_f(\beta; n, k) \frac{\rho_{fs}(\theta_n) \cos^2 \theta_n}{R_f(\beta = 0; n, k) \cos \theta_i \cos \theta_r} S(\theta_n, \beta) + \rho_D + \frac{\rho_V}{\cos \theta_i + \cos \theta_r}, \quad (7)$$

The details and parameters of the model are described in Section B. Based on Equation 7, the Beard Maxwell BRDF function has a total of eight free parameters that can be used to specify the directional reflectance properties of a given material: n , k , ρ_D , ρ_V , σ , ψ , Ω , and B (bias). These terms control the various reflectance characteristics of the material including the amount of back-scatter versus specular-scatter, the amount of shadowing, and the diffuse reflectance of the material.

5.3. Model Inversion Routine and Metrics

Due to the relatively large number of parameters involved in the Beard-Maxwell BRDF model given in Equation 7, we chose to use a two-step approach towards solving for the global optimum of the parameter space. Initial estimates for the parameters are found using an evolutionary algorithm and these initial estimates are then fed into a conjugate-gradient descent routine to find the global optimum solution. Conjugate-gradient descent methods are notoriously sensitive to initial parameter estimates, so this two-step approach was chosen to mitigate the potential of getting stuck in a local optimum solution.

The evolutionary algorithm used to determine initial estimates for each of the BRDF parameters was the Particle Swarm Optimization (PSO) algorithm. In PSO, candidate solutions (particles) are randomly distributed within the parameter space, and allowed to search for the position of their local optimal solution. Each particle's movement is influenced by its local best known position, but is also guided toward the best global positions of all particles. In this way, the trajectory of particles within the parameter space is influenced by both individual search behavior and collective group action (Shi & Eberhart 1998). The implementation of PSO chosen for this work was the PySwarms python module (Miranda et al. 2018). Experimentation revealed that good convergence across randomly seeded runs was achieved when using a 1000 particle swarm with cognitive parameter of 0.5, social parameter of 0.3, and weight parameter of 0.9.

The results of the optimal solution found by the PSO optimizer are used as initial parameter estimates for a conjugate gradient descent routine. The PSO algorithm does not guarantee an optimal solution is found due to the optimizer not relying on the gradient when determining the update velocities of individual particles (Shi & Eberhart 1998). PSO may therefore produce an acceptable solution that is adjacent to global optimal solution, but requires the estimate to be fed into a gradient descent routine in order to determine the true global optimal solution. The gradient descent method chosen is Powell's conjugate gradient descent routine. In Powell's method, the residual function need not be differentiable, as no derivative calculations are used (Powell 1964).

The global optimum solution is defined as the point in the 8 dimensional parameter space $\Theta_{\mathbf{BM}} = (n, k, \rho_D, \rho_V, \sigma, \psi, \Omega, B)$ at which the residual function between the modeled BRDF function and the empirical BRDF measurements is at a minimum. The residual function used in the inversion routine is defined as:

$$P(\Theta_{\mathbf{BM}}) = \sum_{y=0}^{N_m} [\rho_{exp}(\theta_{i_y}, \phi_{i_y}; \theta_{r_y}, \phi_{r_y}) - \rho_{BM}(\Theta_{\mathbf{BM}}; \theta_{i_y}, \phi_{i_y}; \theta_{r_y}, \phi_{r_y})]^2 \quad (8)$$

where N_m is the number of combinations of multi-angle spectral observation angles, ρ_{exp} is the empirical BRDF measurement value with illumination orientation $(\theta_{i_y}, \phi_{i_y})$ and viewing orientation $(\theta_{r_y}, \phi_{r_y})$, and $\rho_{BM}(\Theta_{\mathbf{BM}}; \theta_{i_y}, \phi_{i_y}; \theta_{r_y}, \phi_{r_y})$ is the value of the Beard-Maxwell BRDF function at the view-illumination geometry corresponding to the experimental measurement.

The metric of evaluation utilized to determine the accuracy of the fitting routine is given by the equation of the root mean square error, $RMSE = \sqrt{P(\Theta_{\mathbf{BM}})/N_m}$. This metric is frequently used in studies for fitting empirical BRDF data to theoretical BRDF models (Li & Strahler 1992).

6. RESULTS

6.1. Retrieved BRDF Model Parameters

The derived values for the satellite reflectances used in this study are shown in Table 2. The final global optimum of each of the parameters in the 8 dimensional space are shown, in addition to the $RMSE$ obtained across the different measurements for each stage of the two-step optimization routine.

Validation of BRDF model fitting is difficult due to the abstract nature of certain parameters in Equation 7. One metric that gives us confidence in our results is that the retrieved values of the index of refraction and extinction coefficient match up with spectroscopy studies performed on the solar cells used in silicon solar panels (Green 2008) and aluminum materials (Eismann 2012). It should be kept in mind that our measurement sets are obtained by averaging the BRDF magnitude across the spectral range of $\lambda_L=400$ nm to $\lambda_U=700$ nm for the samples described in Table 1. The index of refraction (n) and extinction coefficient (k) are inherently wavelength-dependent parameters that can describe absorption features on the order of nanometers. The Beard-Maxwell model inversion routines can achieve greater fitting confidence if knowledge of the wavelength dependent reflectance characteristics is used when solving for the global optimum of the wavelength-dependent form of Equation 7 (Westlund & Meyer 2002). This topic should be explored in future studies on the wavelength-dependent reflectance characteristics of spacecraft materials. This remark is especially true given that preliminary empirical measurements of multi-spectral light curves have shown that the multi-spectral information may aid in discerning satellite attitude behaviors from specular reflectance (Alcala & Brown 2009).

We attempted to validate the precision of our inversion routine by performing the optimization procedure outlined in Section 5.3 for 20 randomly initialized simulations. Kernel density estimations of the probability distribution functions (PDFs) of n (top row) and k (bottom row) resulting from these simulations is shown in Figure 5. The values of n and k from the simulation that resulted in the lowest overall RMSE (shown in Table 2) are plotted with dashed red lines on these histograms. From these runs, it is evident that the values in Table 2 were often near the mode of the distributions of n and k . The exception was for the values of n and k retrieved for aluminum 6061 (Figure 5 (a)). The PDF of the retrieved n values appeared to have a bi-modal distribution. The peak of one mode of the distribution was centered near the global-optimum value of n . The global optimum value of k for aluminum 6061, on the other hand, was not adjacent to the mode of the distribution of the simulation results. This variance in the retrieved values of n and k shows that the retrieved results should be taken with care until a broader dataset in terms of both number of wavelengths and number of viewer-illumination angles (N_m) sampled can be obtained. The focus of this study was not on the accuracy of BRDF inversions, but rather on studying if mixtures of BRDFs can contribute to difficulties in training machine learning algorithms to discern spacecraft attitude behaviors. Therefore, the retrieved values of n and k shown in Table 2 were considered sufficient for our purposes.

Renderings of the values from Table 2 for incident light angles of $(\theta_i = 10, \phi_i = 0)$ and $(\theta_i = 50, \phi_i = 0)$ are shown in Figure 6 for the case of the aluminum sample and white paint sample. From these diagrams it is immediately clear that the magnitude and directional properties vary significantly across materials and observer-solar geometry.

For example, when comparing the reflectance properties of white paint (Figure 6 (a)) and aluminum (Figure 6 (b)) under nadir-like illumination conditions of $(\theta_i = 10, \phi_i = 0)$, we can see that the white paint has a diffuse factor that

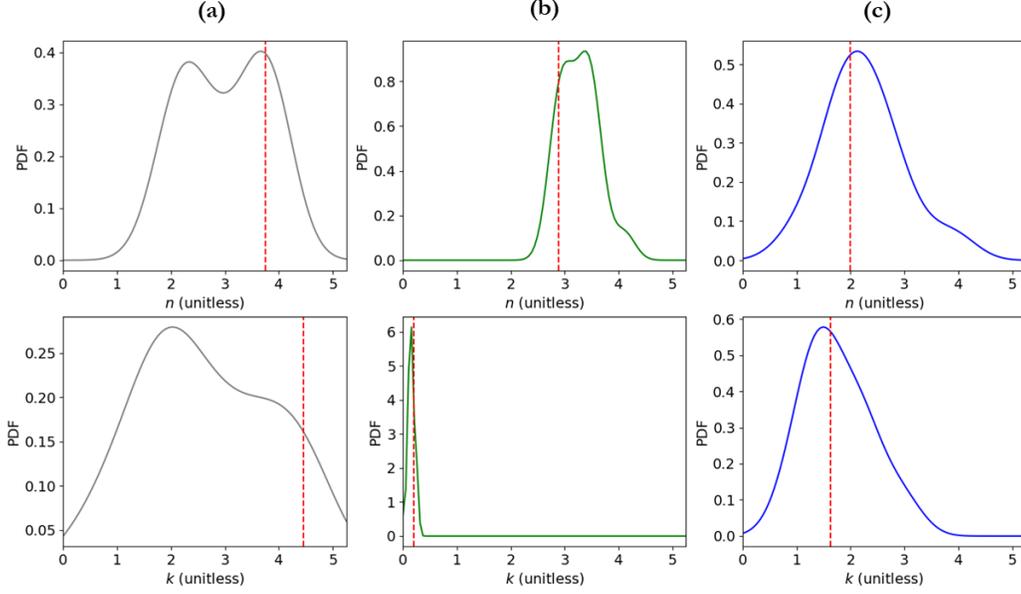


Figure 5. Histogram of the n and k values estimated using a kernel density estimation approach. These values were retrieved after running 20 randomly initialized runs according to the procedure outlined in Section 5.3. The values of the chosen values are marked in red corresponding to the values in Table 2. Histograms are shown for (a) aluminum, (b) solar panels, and (c) glossy white paint.

Table 2. Results of performing PSO followed by Powell’s descent method for the datasets used in this study on the parameters shown in Equation 7 with evaluation metric of $RMSE$.

| Final Parameters | Material | | |
|----------------------------|-------------|----------|--------------|
| | White Paint | Aluminum | Solar Panels |
| n | 1.99 | 3.75 | 2.88 |
| k | 1.63 | 4.47 | 0.21 |
| σ | 0.15 | 0.15 | 0.16 |
| ψ | 0.02 | 0.13 | 0.07 |
| Ω | 0.06 | 0.04 | 0.04 |
| B | 0.06 | 0.01 | 0.003 |
| ρ_D | 0.02 | 0.004 | 0.05 |
| ρ_V | 0.08 | 0.003 | 0.004 |
| RMSE after PSO | 0.016 | 0.007 | 0.05 |
| RMSE after Powell’s Method | 0.01 | 0.001 | 0.004 |

is not present in the aluminum sample as indicated by the increasing slope in BRDF as θ_r increases. In addition, the aluminum sample has a much wider backscatter lobe as indicated by the spike in the diagrams that occurs at the viewing orientation of $(\theta_r = 10, \phi_r = 0)$. A trained machine learning algorithm may construe a back-scattered signal from aluminum as a specular signal due to the back-scatter lobe having a width of approximately 15 degrees. In addition, it should be noted that the overall magnitudes are significantly different with the white paint sample having a peak BRDF of 0.27 and the aluminum having a peak BRDF of 0.04.

Another example of the differences in BRDF properties across materials is seen when comparing the reflectance properties of white paint (Figure 6 (c)) and aluminum (Figure 6 (d)) under oblique illumination conditions of $(\theta_i = 50, \phi_i = 0)$. The diffuse reflectance of white paint is still apparent due to the increasing value of the BRDF with increasing θ_r . In addition, the aluminum sample has a faint back-scatter signal as given by the small increase in BRDF magnitude occurring along the $(\phi_r = 0)$ axis.

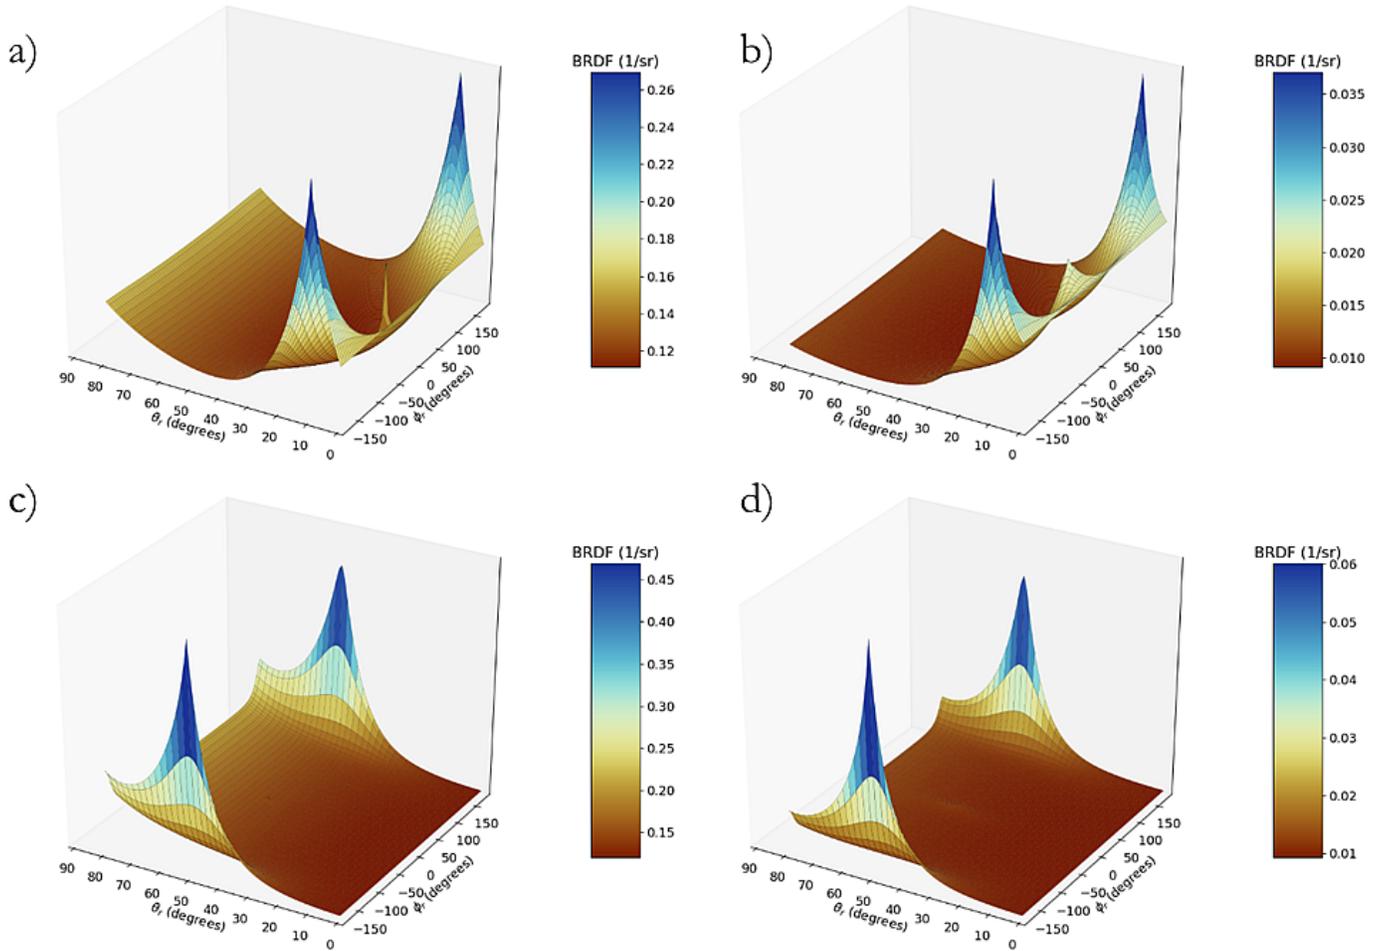


Figure 6. Renderings of the BRDF magnitude as a function of viewing orientation for the white paint sample (left) and the aluminum sample (right). In the top row (a, b) the BRDFs are shown for incident light orientation of $(\theta_i = 10, \phi_i = 0)$. In the bottom row (c, d) the BRDFs are shown for incident light orientation of $(\theta_i = 50, \phi_i = 0)$.

These differences in the BRDF properties of materials should not be ignored, especially for satellite bodies that are non-convex due to protruding wings or antennae. The possibility of observing a specular spike or back-scatter signal can complicate efforts to perform efforts such as satellite attitude estimation or satellite shape estimation.

6.2. Simulated Light Curves

A total of approximately 7000 light curves of varying attitude maneuvers and solar-observer geometries were simulated according to the rotational mechanics and orbital mechanics procedures outlined in Section 4. While an exhaustive description of these light curves is not possible within the confines of this paper, several different light curves are discussed here in order to demonstrate the difficulties a machine learning algorithm might encounter when classifying attitude maneuvers.

Light curves resulting from four different categories of maneuvers (tumbling, stabilizing, accelerating, and inactive) are shown in Figures 7 through 10. In the top row of these figures, the DC signal of the modeled sensor are shown in blue scatter points and the satellite altitude angle (degrees) is shown in red dashed lines, with corresponding y-axes on the left- and right-hand of the plot, respectively. In the bottom row, the proportion of surface areas of the different materials from Table 2 that are visible and illuminated are shown in solid colored lines, and the visual magnitude is shown in dashed red lines. Corresponding y-axes are shown on the left- and right-hand side of the bottom row plots, respectively.

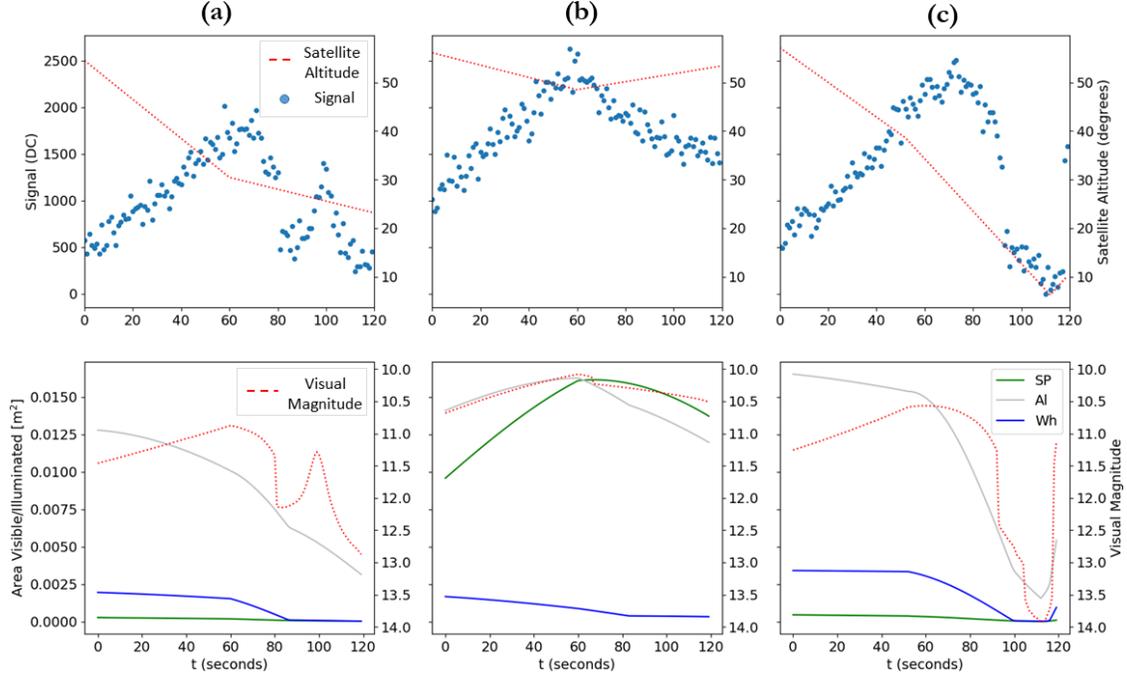


Figure 7. Examples of an inactive satellite. The satellite is considered non-rotating over the entire trajectory of the light curve.

The area that is visible and illuminated to the observer can be written as:

$$A_{v/i}(t) = \sum_j A_j \langle \hat{n}_j^B(t) \cdot \hat{o}^B(t) \rangle \langle \hat{n}_j^B(t) \cdot \hat{s}^B(t) \rangle, \quad (9)$$

Whereas the visual magnitude used in this paper can be written as a summed total of the facet signals as a function of incident illumination and exitant viewing angles (Linares & Furfaro 2016):

$$m_{app}(t) = -26.7 - 2.5 \log_{10} \left(\sum_j \left(\frac{A_j}{R_{RSO/obs}(t)^2} \right) \rho_j(\theta_i(t), \phi_i(t); \theta_r(t), \phi_r(t)) \langle \hat{n}_j^B(t) \cdot \hat{o}^B(t) \rangle \langle \hat{n}_j^B(t) \cdot \hat{s}^B(t) \rangle \right) \quad (10)$$

Note that in comparison to Equation A16, Equation 10 fails to take into account zenith-angle dependent atmospheric blurring and optical blurring effects. This is one of the primary advantages of our simulation efforts versus previous studies on machine learning training using simulated datasets (Linares & Furfaro 2016).

In Figure 7, examples of light curves for non-rotating inactive satellites are shown. For these light curve cases, the oscillations in signal are solely due to the apparent motion of the observer and sun relative to the body facets. Most light curves within this classification can be described by a strong correlation between the signal strength and the satellite altitude, as seen in Figures 7 (a) and (c). An exception to this rule is when a specular signal is observed such as is the case for $t = 100$ seconds in Figure 7 (b). A human operator may classify this light curve as a rotating satellite, which demonstrates the difficulty in discerning the operational status of satellites from light curve composed of mixtures of materials with unique BRDF signatures.

In Figure 8, three different examples of light curves resulting from a tumbling satellite with different angular velocities and axes of rotation are shown. In Figure 8 (a), the satellite is oscillating at a rate of 12 deg/s corresponding to a period of approximately 30 seconds. The full rotation period is easy to discern from the light curve due to the clearly delineated peaks. On the other hand, in Figure 8 (b), the satellite is rotating with an angular velocity of 4 deg/s, but there are few apparent peaks within the time series. Furthermore, the signal strength increases throughout the time series due to a combination of factors: (1) the satellite approaching a closer distance to the earth as time increases, resulting in less pixel smearing according to Equation A12, and (2) an increasing proportion of solar panel becoming visible and illuminated as time increases, as seen in the green line in the bottom row of Figure 8 (b). It is very easy to

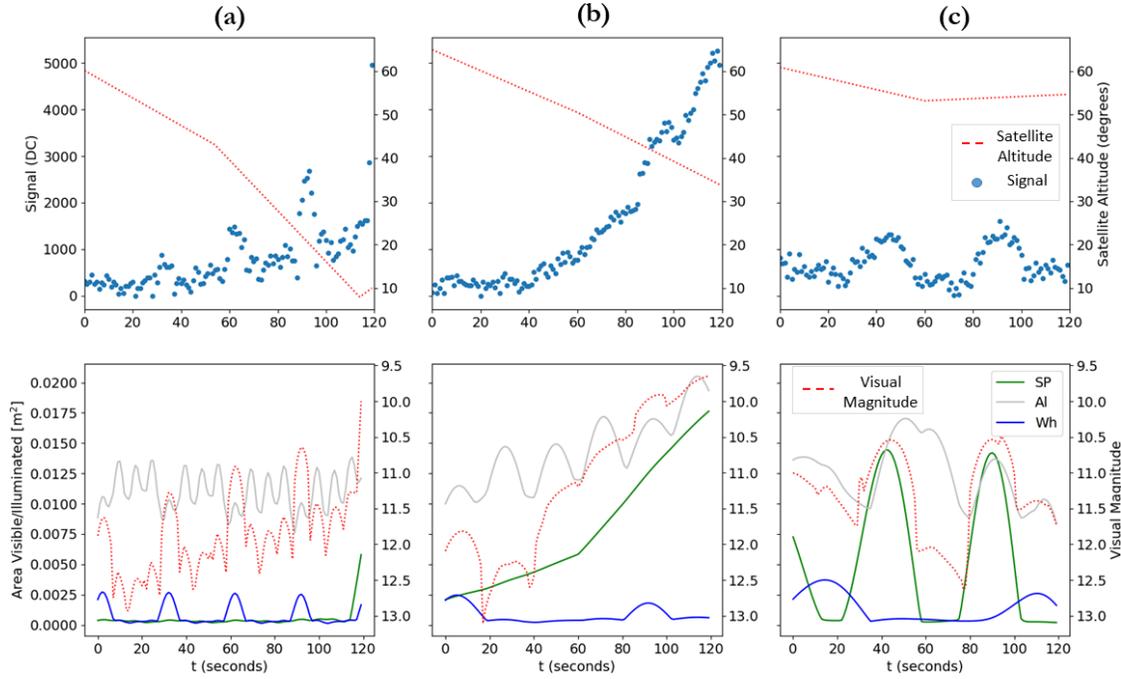


Figure 8. Examples of a tumbling satellite: (a) $\omega_0 = 12$ deg/s about the y axis of the CAD model, (b) $\omega_0 = 4$ deg/s about the y axis, and (c) $\omega_0 = 4$ deg/s about the xz axis.

perceive that this time series could be classified as a stationary satellite according to the light curve examples that are seen in Figure 7. This example in Figure 8 (b) is a good example of how the complex interplay between solar-observer phase angle, distance from satellite to the observer, and mixtures of BRDFs can make discerning satellite attitude status very difficult.

In Figure 9, examples of light curves resulting from a stabilizing satellite are shown. In Figures 9 (a) and (b), the satellite is rotating with a slow angular velocity of 4 deg/s. The rotation behavior of the RSO is obscured by noise and blurring of the signal, making it possible that these light curves could be classified as resulting from inactive satellites. On the other hand, in Figure 9 (c), there are many clearly defined peaks in the light curve due to highly specular materials such as solar panels (green line) and white paint (blue line) becoming visible throughout the time series. It is easy to discern that this satellite is rotating, but difficult to perceive that it is accelerating due to noise within the signal.

Finally, in Figure 10 we show several examples of an oscillating stabilizing satellite. In Figures 10 (a) and (c), it is easy to make out satellite rotation behaviors due to the strong specular and back-scatter signals resulting from both solar panels (green line) and white paint (blue line). While evident that it is rotating, noise within the signal makes it unclear from an operator perspective whether it is accelerating, stabilizing, or maintaining rotation rate. In Figure 10 (b), a light curve due to a stabilizing satellite is shown where the DC signal plateaus at the same time period as the satellite stabilizes. Prior to this plateau, the signal is characterized by varying signal about a smooth trend-line. This light curve behavior could potentially be learned by a machine learning algorithm in order to detect stabilizing or accelerating maneuver behaviors.

6.3. CNN Results

6.3.1. Retrieved CNN Architecture

The results of testing our trained CNN approach against our simulated data are presented in this section. Of the approximately 7500 light curves that were generated over the four different maneuvers, 80% of each class were used as a training dataset while the remainder of each class were used for testing the performance of our algorithm.

Keras and Tensorflow were used in order to train and test our neural network architecture in this work (Chollet et al. 2015). Using the various network layers outlined in Section 2.2 and the training metrics such as early stopping outlined in Section 2.3, an optimal neural network architecture was derived by carrying out a process of performing hyper-parameter tuning. Various hyper-parameters that were tuned in order to achieve the best possible model performance

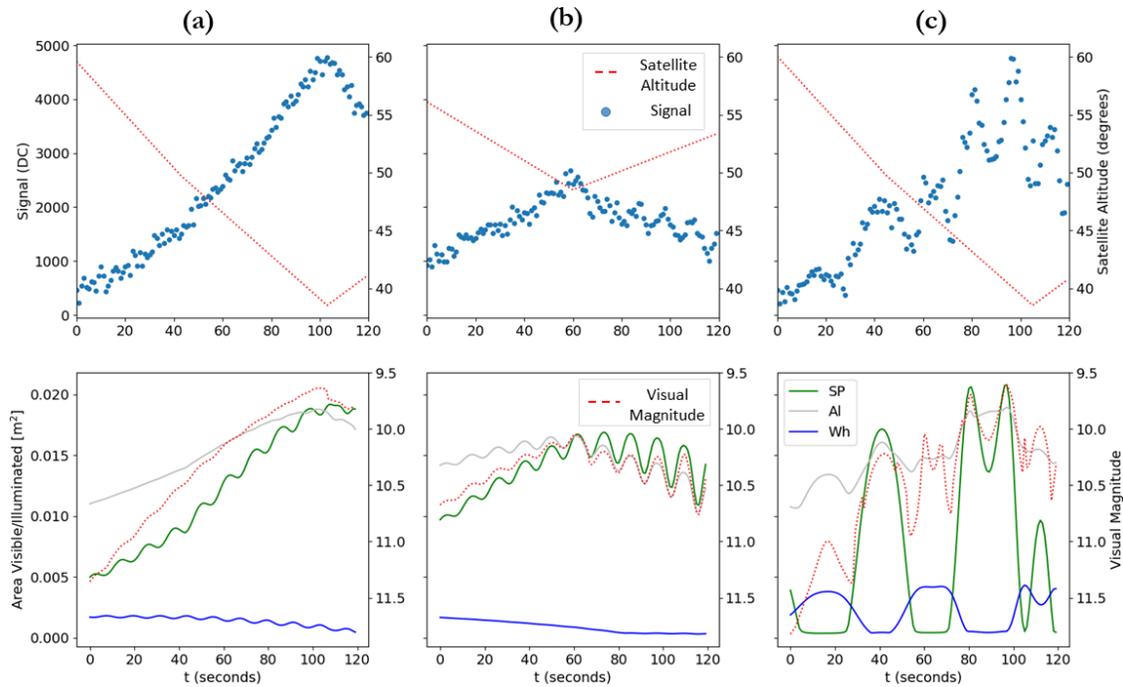


Figure 9. Examples of an oscillating accelerating satellite: (a) $T = 12$ seconds, $\omega_0 = 6$ deg/s, and $\gamma = 45$ seconds about the xz -axis of the CAD model, (b) $T = 16$ seconds, $\omega_0 = 6$ deg/s, and $\gamma = 45$ seconds about the z -axis of the CAD model, and (c) $T = 48$ seconds, $\omega_0 = 24$ deg/s, and $\gamma = 45$ seconds about the x -axis of the CAD model.

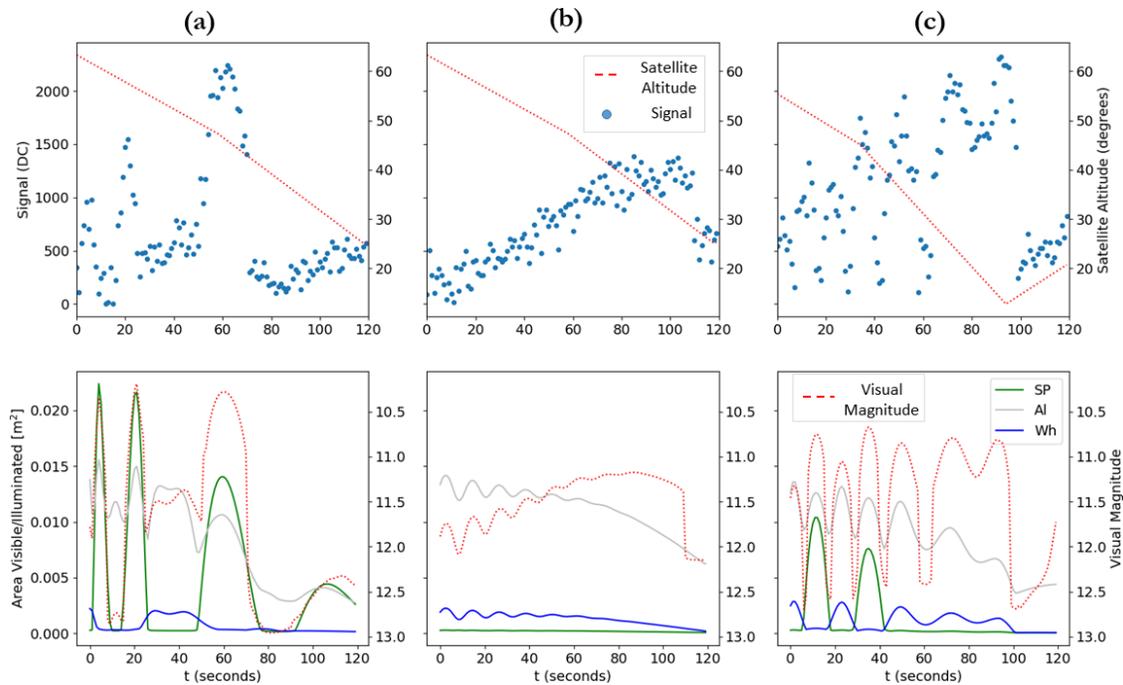


Figure 10. Examples of light curves resulting from an oscillating stabilizing satellite: (a) $T = 48$ seconds, $\omega_0 = 24$ deg/s, and $\gamma = 120$ seconds about the xz -axis, (b) $T = 12$ seconds, $\omega_0 = 6$ deg/s, and $\gamma = 120$ seconds about the x -axis, and (c) $T = 24$ seconds, $\omega_0 = 12$ deg/s, and $\gamma = 45$ seconds about the x -axis.

for our dataset, including: the learning rate of the Adam optimizer, the number of convolution layers, the convolution kernels used at each layer, the size of the fully connected layers, and the dropout rates of the fully connected layers. A batch size of 32 light curves was used in order to train the CNN.

The derived neural network architecture as a result of the hyper-parameter tuning process is shown in Figure 2. As can be seen, our study employed a five layer structure. The model consists of three convolution layers followed by two fully connected layers. The three convolution layers use 5, 10, and 10 size convolution kernels, respectively. They also consist of 6, 24 and 24 features per layer, respectively. After each convolution layer, a procedure of performing ReLU activation, batch-normalization, and max-pooling (4x1 pooling kernel) is carried out.

The fully connected layers are of size 120 and are always followed by a ReLU activation function, except for the output layer. The dropout rate used in the process of training the fully connected layers is 0.6 for both of the fully-connected layers.

6.3.2. CNN Performance

Our model achieved a peak overall classification accuracy of 86.2% on our testing dataset before the training process was stopped in order to reduce the potential for over-fitting. A confusion matrix representing the performance of our CNN across the different classes of satellite maneuvers is shown in Figure 11. Note that this confusion matrix is derived for a random sampling of light curves to be used in the testing and training data-sets but that results showed that this performance held to within $\pm 7.5\%$ across 10 different random drawings of the light curves to be utilized in the training and testing data-sets.

From the results in Figure 11, it is clear that our CNN performed well for the case of tumbling and stabilizing satellites. We were able to accurately classify 94% of both tumbling and stabilizing satellites across a broad range of observation and illumination geometries, with the most common false-alarm class among these two maneuver classes being the accelerating maneuver category. This result is promising because it means that our system has a strong ability to correctly classify an out-of-control RSO.

Our CNN struggled with the task of classifying accelerating satellite maneuvers with a classification accuracy of 78%. Tumbling and stabilizing maneuvers were common false-alarms with 6% of the testing data-set being labeled as these maneuvers. This outcome is understandable given that these maneuvers represent different spin-states of the satellite and that some observed light curves occurred when the satellite was still tumbling with a high angular velocity and/or sharp transitions in solar-observer phase angle. More concerning is the false alarm rate of 10% for acceleration maneuvers being classified as inactive satellites; this could be due to the difficulties with properly classifying inactive RSOs that exhibit specular reflections.

Our CNN performed poorly for the case of classifying an inactive satellite maneuver given by the low accuracy of just 60%. The CNN had a relatively high probability of classifying several light curves representing inactive satellites as either accelerating (22%) or stabilizing (17%). Further investigation of the mis-classified light curves representing inactive RSO attitude states revealed that this behavior could likely be attributed to two major factors: glints in the light curves and system level noise.

Figure 12 shows the classification accuracy for the training dataset of inactive RSO maneuvers across two different histograms of the metadata. In Figure 12 (a), the classification accuracy is broken down across bins of the peak SNR of the light curve as calculated according to Equation A15. The classification accuracy when the peak SNR in the light curve is less than 15 is only 47%. However, the classification accuracy steadily rises to 71% for light curves with a peak SNR level of ≥ 30 . This result indicates that system level noise due to factors such as atmospheric turbulence blurring and background signal noise can complicate training of a CNN for RSO attitude prediction. This means that our system's predictions should be evaluated more harshly in scenarios where the RSO is at a low altitude, leading to more blurring due to atmospheric turbulence.

In Figure 12 (b), the classification accuracy for inactive RSO maneuvers is broken down across bins of the number of peaks within the visual magnitude of the signal as given by Equation 10. In the case of a visual magnitude plot with ≤ 1 peak, the light curve was assumed to have no specular or back-scatter glints within the signature such as in Figure 7 (b). For the case of a visual magnitude plot with ≥ 2 peaks, it was assumed that one or more glints occurred within the light curve such as in Figure 7 (a). The classification accuracy was marginally higher for the case of light curves without specular glints (65%) when compared to light curves with one or more glints (51%).

Our CNN clearly struggled with discerning system level noise and specular glints versus attitude maneuvers of the RSO. Our results are bolstered by a study performed by Dianetti & Crassidis (2018), in which they attempted to

| | | Predicted | | | |
|-------|--------------|--------------|----------|-------------|----------|
| | | Accelerating | Inactive | Stabilizing | Tumbling |
| Truth | Accelerating | 0.78 | 0.10 | 0.06 | 0.06 |
| | Inactive | 0.22 | 0.60 | 0.17 | 0.01 |
| | Stabilizing | 0.02 | 0.03 | 0.94 | 0.01 |
| | Tumbling | 0.05 | 0.01 | 0.00 | 0.94 |

Figure 11. A confusion matrix representing the classification performance of our CNN across the different classes used in this study.

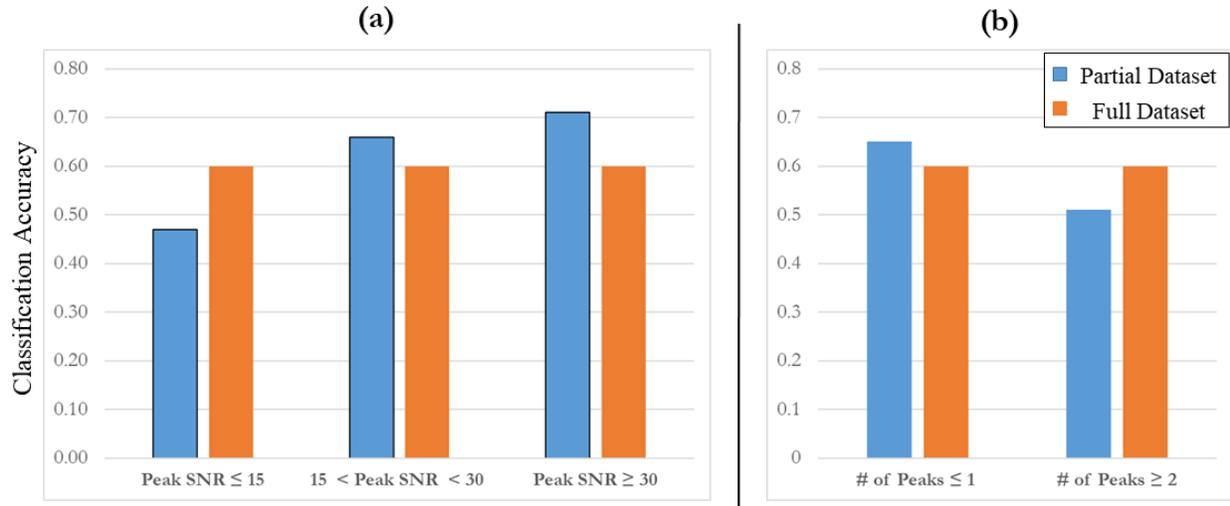


Figure 12. Bar charts representing the classification accuracy of the trained CNN when classifying inactive RSO light curves across (a) varying levels of peak SNR, and (b) varying numbers of peaks in the observed light curve. The orange bars represent the classification accuracy across all inactive RSO light curves in the testing data-set.

use frequency domain decomposition of light curves to understand rotation states of single-BRDF RSOs. Diannetti showed that satellite attitude state can be difficult to predict when there are low SNR conditions and glint signatures present within the observed light curve. Clearly, a CNN struggled with these same exact issues.

7. CONCLUSION

In this study, CNNs were trained for the task of inferring the operational status of RSOs from simulated light curves. The CNN was trained on light curves resulting from the Georgia Tech RANGE-A satellite performing four different maneuvers: tumbling, accelerating in rotational rate, stabilizing in rotational rate, and inactive (or stable in rotation rate). The CNN achieved an overall classification accuracy of 86.2% across the four maneuver classes. A confusion matrix analysis of the different classes of maneuvers suggested that our model performed best when classifying tumbling and accelerating RSOs (94% accuracy) and worst at classifying inactive RSOs (60% accuracy). Further investigation showed that the reduction in accuracy when classifying inactive RSOs could be attributed to BRDF-induced specular reflections being misinterpreted as maneuvers, and system-level noise being mistaken for a varying rotational rate. These results suggest that CNNs have strong potential for aiding in the problem of classifying satellite attitude status from light curves, but that machine learning research must focus on developing training sets and pre-processing techniques that account for these complications.

There are many approaches that can be investigated in order to improve our classification performance in future studies. One approach is generating a more robust dataset that can be used to train our model with an larger number of examples of inactive attitude maneuvers across different observation geometries. Another approach is to consider using multi-spectral light curves for training and testing our CNN system. For example, [Alcala & Brown \(2009\)](#) have shown that multi-spectral light curves can potentially aid in differentiating between BRDF signatures and satellite rotation status. Another approach is to use different machine learning techniques that have seen success in the field of

SDA such as random forests (Dao et al. 2019) and different CNN architectures (McQuaid et al. 2018). These are all active areas of research that will be elaborated upon in future studies.

8. ACKNOWLEDGMENTS

This work has been funded by Georgia Tech Research Institute Internal Research and Development (IRAD) funding.

APPENDIX

A. LIGHT CURVE SIMULATION

A.1. Propagation of Signal from RSO in Orbit to Telescope Aperture

The total irradiance, M_{orbit} , in units of $Watts/m^2$ onto a facet of an RSO in orbit can be obtained by modeling the sun as an idealized blackbody and assuming that the sun is an isotropic point source (Budding & Demircan 2007):

$$M_{orbit} = \int_{\lambda_L}^{\lambda_U} \left(\frac{r_{sun}^2}{(1 AU)^2} \right) \frac{2\pi hc^2}{\lambda^5 \left[\exp\left(\frac{hc}{\lambda K_B T_{sun}}\right) - 1 \right]} d\lambda, \quad (A1)$$

where T_{sun} is the temperature of the sun, K_B is Boltzmann's constant, c is the speed of light, λ is the wavelength of interest, h is Planck's constant, $1 AU$ is the mean earth-sun distance and r_{sun} is the radius of the sun. The wavelength-dependent irradiance was integrated over the upper (λ_U) and lower transmission (λ_L) bounds of the spectral filter of interest to give the wavelength-independent irradiance in this equation.

We assume that the distance between facets along the RSO is sufficiently small relative to the RSO-to-sun distance, $R_{RSO/sun}(t)$, that M_{orbit} can be considered constant over all RSO facets. The radiance, L_{orbit_j} , leaving the j^{th} facet of the RSO body into a given observer direction $(\theta_r(t), \phi_r(t))$ can be written according to:

$$L_{orbit_j}(t, \theta_r(t), \phi_r(t)) = M_{orbit} \rho_j(\theta_i(t), \phi_i(t); \theta_r(t), \phi_r(t)) \langle \hat{n}_j^B(t) \cdot \hat{s}^B(t) \rangle, \quad (A2)$$

where $(\theta_i(t), \phi_i(t))$ is the time-dependent direction of the sun relative to the facet in the body-frame, \hat{n}_j^B is the unit normal vector denoting the orientation of the j^{th} facet relative to the body frame, \hat{s}^B is the unit normal vector specifying the solar orientation in the body frame, and ρ_j is the observer/illumination direction dependent BRDF of the j^{th} facet's material. Note that the BRDF of spacecraft materials has been shown to be highly dependent on wavelength (Bédard et al. 2015; Reyes & Cone 2018), but is taken to be averaged over the range of the spectral filter $[\lambda_L, \lambda_U]$ in this study. The observer-illuminator geometry parameters $(\theta_i(t), \phi_i(t); \theta_r(t), \phi_r(t))$ are all time-dependent due to the attitude and orbit of the RSO, but for simplicity, the time dependence will be dropped from this point forward.

The average irradiance onto an observing telescope aperture due to the j^{th} facet, E_{ap_j} , in units of $Watts/m^2$ can then be written by taking into account the solid angle subtended by the RSO:

$$E_{ap_j}(t; \theta_i, \phi_i; \theta_r, \phi_r) = L_{orbit_j}(t, \theta_r, \phi_r) \tau_{atm}(\theta_r) \left(\frac{A_j}{R_{RSO/obs}(t)^2} \right) \langle \hat{n}_j^B(t) \cdot \hat{o}^B(t) \rangle, \quad (A3)$$

where \hat{o}^B is the unit normal vector specifying the observer orientation in the body frame, A_j is the surface area of the j^{th} facet, $\tau_{atm}(\theta_r)$ is the zenith-angle dependent atmospheric transmittance, and $R_{RSO/obs}(t)$ is the time dependent distance of the RSO to the observer telescope aperture.

In order to obtain the total expected irradiance, E_{ap} , onto the aperture, the signal due to all facets is summed:

$$E_{ap}(t; \theta_i, \phi_i; \theta_r, \phi_r) = \sum_j E_{ap_j}(t; \theta_i, \phi_i; \theta_r, \phi_r) \quad (A4)$$

The flux due to a point source that reaches the focal plane, Φ_{FPA} , in units of $[Watts]$ can be written according to:

$$\Phi_{FPA}(t) = E_{ap}(t; \theta_i, \phi_i; \theta_r, \phi_r) \tau_{opt}(\bar{\lambda}) \left(\frac{\pi D_{ap}^2}{4} \right) \quad (A5)$$

where $\tau_{opt}(\bar{\lambda})$ is the average optical transmittance of the filter over the bandpass of interest, D_{ap} is the aperture diameter, and $\bar{\lambda} = (\lambda_L + \lambda_U)/2$ is the average wavelength under consideration.

The signal can then be converted to units of e^- /second by taking into account the wavelength of bandpass of interest and the quantum efficiency of the focal plane array electronics (Shell 2010):

$$q_{S0}(t) = \Phi_{FPA}(t) \left(\frac{\bar{\lambda}}{hc} \right) QE(\bar{\lambda}) \quad (A6)$$

A.2. Propagation of Energy From Sky Background Onto Focal Plane Array

By assuming that the sky background is an extended source, we can derive the spectral irradiance onto a pixel along the focal plane array according to the following equation for a telescope of $f_{\#} = f/D_{ap}$, where f is the focal length of the telescope (Coder & Holzinger 2016):

$$E_{sky} = \frac{L_{sky} \pi \tau_{optics}(\bar{\lambda})}{(1 + 4f_{\#}^2)} \quad (A7)$$

where L_{sky} is the photon radiance at the telescope aperture due to background sky pollution in units of photons/s/m²/sr (Coder & Holzinger 2016):

$$L_{sky} = M_{orbit} \left(\frac{\bar{\lambda}}{hc} \right) \left(\frac{180}{\pi} \right)^2 3600^2 10^{-0.4(I_{sky} + m_{v_0})} \quad (A8)$$

In the above equation $m_{v_0} = 26.4$ is the magnitude of the Vega star which serves as the magnitude-zero source (Budding & Demircan 2007), and I_{sky} is the local background sky radiant intensity due to light pollution, moonlight, and starlight in units of visual magnitude per arcsec².

A final expression for the photon flux per pixel due to sky background, $q_{p,sky}$, in units of e^- /pixel/second can be written as (Shell 2010):

$$q_{p,sky} = \frac{L_{sky} \pi \tau_{optics}(\bar{\lambda}) QE(\bar{\lambda}) p^2}{(1 + 4f_{\#}^2)}, \quad (A9)$$

where p is the pixel pitch.

A.3. Energy Loss Due to Atmospheric Turbulence Blur and Sampling

Inhomogeneities in the temperature of air within the atmosphere cause refractive index variations over the path from the surface of the earth to space. This refractive index gradient along the propagation path causes light to be received in a large variety of angles of incidence at the optical aperture (Zamek & Yitzhaky 2006). Over short integration time scales (on the order of a few milliseconds), propagation through a turbulent atmosphere leads to phenomena such as image dancing, while over long time scales (on the order of hundreds of milliseconds) many dancing resolved images are collected which leads to image blurring (Fried 1966). Long-term blurring effects on imagery collected from optical telescopes arise primarily due to the effects of atmospheric turbulence on the phase rather than the amplitude of the propagating wavefront (Lawrence & Strohbehm 1970). Consequently, in an imaging context, it can be assumed that energy is conserved when examining the effects of atmospheric turbulence on the received irradiance at the telescope aperture.

A plane wave approximation is met when the path length difference from an RSO in orbit to points along the extent of the observer aperture is negligible compared to the wavelength (Zamek & Yitzhaky 2006):

$$D_{ap} \ll \sqrt{\bar{\lambda} R_{RSO/obs}(t)} \quad (A10)$$

For the observer-to-RSO distances and optical wavelengths considered in this paper, Equation A10 is met and a plane wave treatment is used.

For a plane wave, the full-width half maximum (FWHM) of the point spread function (PSF) due to atmospheric seeing is approximately parameterized by the wavelength-dependent Fried parameter of the atmosphere at a wavelength of 500 nm, $r_0(\theta_r)$: $FWHM_{atm} = \bar{\lambda}/r_0(\theta_r)$ (Fried 1966). The value of r_0 is non-linearly dependent on the observation

zenith angle, θ_r . The equation for the dependence of the Fried parameter on zenith angle for a plane wave under the assumption of a constant atmospheric turbulence profile is given by (Fried 1966):

$$r_0(\theta_r) = 2.1 \left[1.46 \sec(\theta_r) \left(\frac{2\pi}{500 \text{ nm}} \right)^2 C_n^2 \right]^{-3/5} \quad (\text{A11})$$

where C_n^2 is the altitude-independent refractive index structure constant. Equation A11 has been empirically shown to be valid for observing zenith angles up to 45 degrees (Lawrence & Strohbehm 1970).

For a moving RSO, the number of pixels occupied by the RSO along the sensor focal plane, $m(t)$, grows according to:

$$m(t) = m_0(t) + \frac{\sqrt{m_0(\lambda, t)} \omega t_{int}}{IFOV}, \quad (\text{A12})$$

where IFOV is the instantaneous field of view of a pixel on the sensor focal plane in units of *radians/pixel*, $m_0(\lambda, t)$ is the number of pixels initially subtended by the RSO, t_{int} is the integration time of the telescope system, and ω is the angular velocity of the RSO in units of *rad/sec* under the assumption of a telescope operating in stare mode (Coder & Holzinger 2016). The IFOV of the telescope is defined by $IFOV = 2 \arctan(p/2f)$, where p is the pixel pitch, and f is the lens focal length.

The wavelength dependent number of pixels occupied by an RSO of projected area A_{SO} at a distance $R_{RSO/obs}(t)$ to the telescope can then be approximately defined as (Coder & Holzinger 2016):

$$m_0(t) = \left(\frac{A_{SO}}{R_{RSO/obs}(t) \times IFOV} + \frac{1}{IFOV} \left(\frac{\bar{\lambda}}{r_0(\theta_r)} \right) \right)^2 \quad (\text{A13})$$

Equation A13 gives the number of pixels across which the signal due to the RSO will be smeared. However, algorithms for calculating light curves will normally only integrate over a small rectangular transect that is aligned along the streak of the light curve in a collected image. In this paper, we assume that the rectangular extent has a transverse side length of Δl . The total number of pixels over which a light curve extraction algorithm will integrate is then given by:

$$m_{\Delta l} = \Delta l (m_0(t)) \quad (\text{A14})$$

The fractional amount of energy lost due to the sampling algorithm of a streak detector can then be approximately calculated by $m_{\Delta l}/m(\lambda, t)$.

In this study, the parameter r_0 is defined for the site of interest at zenith angle $\theta_r = 0$ degrees using measures of atmospheric seeing obtained from previous studies (Coder & Holzinger 2013; Lawrence & Strohbehm 1970). Equation A11 is then inverted to determine C_n^2 and the resultant value is used to infer the value of $r_0(\theta_r)$ for the RSO geometries over the light curve time series.

A.4. Noise Terms

The instantaneous Signal-to-Noise Ratio (SNR) of an RSO observed by a telescope on the ground can be derived by modeling a number of noise terms, including (1) the photon arrival process on the telescope's CCD as a Poisson process, (2) dark current of sensor electronics, and (3) read noise in the process of conversion from analog to digital units. By assuming a constant background noise spectral signature, the instantaneous SNR can be written as (Schildknecht 1994):

$$SNR(t) = \frac{q_{SO}(t)t_{int}}{\sqrt{q_{SO}(t)t_{int} + 2m(t)[q_{p,dark} + q_{p,sky}]t_{int} + q_{read}/n^2}}, \quad (\text{A15})$$

where $q_{p,dark}$ is the dark noise of the CCD [$e^-/pixel/s$], $q_{p,sky}(\lambda)$ is the wavelength-dependent sky signature reaching a pixel of the focal plane [e^-/s], q_{read} is the read noise of the sensor electronics [$e^-/pixel/s$], t_{int} is the sensor integration time, $m(\lambda)$ is the wavelength-dependent number of pixels occupied by the RSO along the focal plane, and n is a pixel binning factor.

The mean and variance of a Poisson distribution representing the photon arrival rate are equal to the rate parameter of the distribution (Coder & Holzinger 2016). Using this representation in combination with the assumption of

independence of the random variable noise terms in Equation A15, the mean and variance of the received light curve signal, $s_r(t)$, in units of electrons can be written as a Gaussian distribution following (Schildknecht 1994):

$$\begin{aligned}\mu_r(t) &= \left(\frac{m_{\Delta l}}{m(t)}\right) q_{SO} t_{int} \\ \sigma_r^2(t) &= \left(\frac{m_{\Delta l}}{m(t)}\right) q_{SO} t_{int} + m_{\Delta l} [q_{p,dark} + q_{p,sky}] t_{int} + q_{read}/n^2,\end{aligned}\tag{A16}$$

where the term $m_{\Delta l}/m(t)$ is meant to signify the loss of a percentage of RSO signal due to the sampling of the light streak algorithm. Using Equation A16, the instantaneous signal received due to the RSO for a given light track is derived as a random variable that is meant to simulate the combined effects of motion blur, optical blur, and noise due to the electronics of the system.

This term can be converted into Digital Counts (DC) by using knowledge of the full-well depth of sensor in units of electrons, N_{well} , and of the dynamic range of the system in units of bits, $N_{DR} = 2^b$. Under this assumption, the received signal in units of electrons can be converted into digital counts according to:

$$DC_r(t) = \mathbf{int} \left(\frac{s_r(t) N_{DR}}{N_{well}} \right)\tag{A17}$$

where the integer operator truncates the signal and further contributes to the noise of the received signal.

B. BEARD MAXWELL BRDF MODEL PARAMETERS

In Equation 7, the incident and exitant polar angles are specified by the time series of solar, observer, and macro-facet orientations within the body frame:

$$\begin{aligned}\theta_i &= \langle \hat{n}_j^B(t) \cdot \hat{s}^B(t) \rangle \\ \theta_r &= \langle \hat{n}_j^B(t) \cdot \hat{o}^B(t) \rangle,\end{aligned}\tag{B18}$$

The half-angle between the incident and exitant directions (θ_n) and the phase angle of the solar-observer geometry (β) are written respectively as:

$$\theta_n = \cos^{-1} \left(\frac{\cos \theta_i + \cos \theta_r}{2 \cos \beta} \right)\tag{B19}$$

$$\beta = \frac{1}{2} \cos^{-1} [\cos \theta_i \cos \theta_r + \sin \theta_i \sin \theta_r \cos(\phi_r - \phi_i)]\tag{B20}$$

The term $R_f(\beta; n, k)$ denotes the effective Fresnel reflectance of a material with complex index of refraction of $n^* = n + ik$, perpendicular reflectance coefficient of r_s , and perpendicular reflectance coefficient of r_p . The values of n and k capture the wavelength dependency of first surface reflection in terms of the index of refraction and extinction coefficients of the material (Maxwell et al. 1973). In this study, the values of the Beard-Maxwell model are fit over a broad bandpass such that the values of n and k are considered to be averaged over wavelength:

$$\begin{aligned}R_f(\beta; n, k) &= \frac{1}{2} (|r_s(n^*)|^2 + |r_p(n^*)|^2) \\ r_s(n^*) &= \frac{\cos \beta - \sqrt{(n^*)^2 - \sin^2 \beta}}{\cos \beta \sqrt{(n^*)^2 - \sin^2 \beta}} \\ r_p(n^*) &= \frac{(n^*)^2 \cos \beta - \sqrt{(n^*)^2 - \sin^2 \beta}}{(n^*)^2 \cos \beta + \sqrt{(n^*)^2 - \sin^2 \beta}},\end{aligned}\tag{B21}$$

The term $\rho_{fs}(\theta_n)$ denotes the first-surface reflectance. This term represents light being reflected in the specular direction off of a collection of micro-facets determined by a Cauchy distribution that make up the texture of the j^{th} macro-facet pointing in direction $\hat{n}_j^B(t)$ (Westlund & Meyer 2002; Montanaro 2007):

$$\rho_{fs}(\theta_n) = \frac{R_f(\beta = 0; n, k) B}{4 \cos^3 \theta_n [\sigma^2 + \tan^2 \theta_n]},\tag{B22}$$

where σ is the mean square value of the micro-facet slope normal directions of the distribution, and B is a scaling parameter meant to fit the overall magnitude of the specular reflectance.

The shadowing and obscuration term, $S(\theta_n, \beta)$, accounts for the height distribution of the micro-facets. This function was derived according to fitting free parameters using empirical data (Maxwell et al. 1973). This function, can be written as (Montanaro 2007):

$$S(\theta_n, \beta) = \frac{1 + \frac{\theta_n}{\Omega} \exp^{-2\beta/\psi}}{1 + \frac{\theta_n}{\Omega}} \quad (\text{B23})$$

where the free parameters Ω and ψ modify the falloff of the shadowing function in the forward-scattered and back-scattered directions, respectively.

Finally, the terms ρ_D and ρ_V represent the diffuse and directional volumetric reflectance of the material, respectively. Note that in the original Beard-Maxwell model, only one volumetric component was utilized (Maxwell et al. 1973). In the enhanced version of the model used in this study both the Lambertian reflectance, ρ_D , and the directional diffuse reflectance, ρ_V , are used as fitting parameters (Westlund & Meyer 2002).

REFERENCES

- Alcala, C. M., & Brown, J. H. 2009, Space Object Characterization Using Time-Frequency Analysis of Multi-spectral Measurements from the Magdalena Ridge Observatory, Tech. rep., Air Force Research Lab Space Vehicles Directorate
- Badura, G., Valenta, C. R., Gunter, B., Renegar, L., & Wu, D. 2019, *amos*, 27
- Bedard, D., Lévesque, M., & Wallace, B. 2011, in Proc. of the Advanced Maui Optical and Space Surveillance Technologies Conf, 1–10
- Bédard, D., Wade, G. A., & Abercromby, K. 2015, *Journal of Spacecraft and Rockets*, 52, 1038
- Bradley, B. K., & Axelrad, P. 2014, in Univ. of Colorado. International Space Symposium for Flight Dynamics
- Budding, E., & Demircan, O. 2007, Introduction to astronomical photometry, Vol. 6 (Cambridge University Press)
- Caruana, R., Lawrence, S., & Giles, C. L. 2001, in Advances in neural information processing systems, 402–408
- Chollet, F., et al. 2015, Keras, <https://keras.io>
- Coder, R., & Holzinger, M. 2013, in Advanced Maui Optical and Space Surveillance Technologies Conference
- Coder, R. D., & Holzinger, M. J. 2016, *Acta Astronautica*, 128, 669
- Cornell. 2002, Reflectance Data, Cornell University Program of Computer Graphics, <https://www.graphics.cornell.edu/online/measurements/reflectance/index.html>
- Dai, J. S. 2015, *Mechanism and Machine Theory*, 92, 144
- Dao, P., Haynes, K., Gregory, S., et al. 2019
- Dianetti, A. D., & Crassidis, J. L. 2018, in 2018 AIAA Guidance, Navigation, and Control Conference, 1605
- DiBona, P., Foster, J., Falcone, A., & Czajkowski, M. 2019, in Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)
- Eismann, M. T. 2012, Hyperspectral remote sensing
- Fan, S., Friedman, A., & Frueh, C. 2019, *amos*, 23
- Fried, D. L. 1966, *JOSA*, 56, 1372
- Früh, C., Kelecy, T. M., & Jah, M. K. 2013, *Celestial Mechanics and Dynamical Astronomy*, 117, 385
- Fulcoly, D. O., Kalamaroff, K. I., & Chun, F. 2012, *Journal of Spacecraft and Rockets*, 49, 76
- Furfaro, R., Linares, R., & Reddy, V. 2018, in AMOS Technologies Conference, Maui Economic Development Board
- Furfaro, R., Linares, R., & Reddy, V. 2019, in AMOS Technologies Conference, Maui Economic Development Board, Kihei, Maui, HI
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep learning (MIT press)
- Green, M. A. 2008, *Solar Energy Materials and Solar Cells*, 92, 1305
- Gunter, B. C., Davis, B., Lightsey, G., & Braun, R. D. 2016, Proceedings of the AIAA/USU Conference on Small Satellites, Session V: Guidance and Control. <http://digitalcommons.usu.edu/smallsat/2016/S5GuidCont/3/>
- Hall, D., Calef, B., Knox, K., Bolden, M., & Kervin, P. 2007, in The 2007 AMOS Technical Conference Proceedings, Maui Economic Development Board, Inc. Kihei, Maui, HI, 464–475
- Haselsteiner, E., & Pfurtscheller, G. 2000, *IEEE transactions on rehabilitation engineering*, 8, 457
- Holzinger, M., & Jah, M. 2018, Challenges and Potential in Space Domain Awareness, American Institute of Aeronautics and Astronautics
- Hou, Q., Wang, Z., Su, J., & Tan, F. 2019, *Coatings*, 9, 193
- Kingma, D. P., & Ba, J. 2014, arXiv preprint arXiv:1412.6980

- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, in *Advances in neural information processing systems*, 1097–1105
- Lawrence, R. S., & Strohbehn, J. W. 1970, *Proceedings of the IEEE*, 58, 1523
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proceedings of the IEEE*, 86, 2278
- Li, X., Chen, S., Hu, X., & Yang, J. 2019, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2682–2690
- Li, X., & Strahler, A. H. 1992, *IEEE transactions on Geoscience and Remote Sensing*, 30, 276
- Linares, R., & Furfaro, R. 2016, in *2016 19th International Conference on Information Fusion (FUSION)*, IEEE, 1140–1146
- Marana, A. N., Velastin, S., Costa, L., & Lotufo, R. 1997
- Marschner, S. R., Westin, S. H., Lafortune, E. P., & Torrance, K. E. 2000, *Applied Optics*, 39, 2592
- Maxwell, J., Beard, J., Weiner, S., Ladd, D., & Ladd, S. 1973, *Bidirectional Reflectance Model Validation and Utilization.*, Tech. rep., Environmental Research Institute of Michigan Ann Arbor Infrared and Optics Division
- McQuaid, I., Merkle, L. D., Borghetti, B., Cobb, R., & Fletcher, J. 2018, in *The Advanced Maui Optical and Space Surveillance Technologies Conference*
- Miranda, L. J. V., et al. 2018, *J. Open Source Software*, 3, 433
- Montanaro, M. 2007, Rochester Institute of Technology, *DIRS Technical Report 2007-83*, 174
- Parisi, R., Di Claudio, E. D., Lucarelli, G., & Orlandi, G. 1998, in *ISCAS'98. Proceedings of the 1998 IEEE International Symposium on Circuits and Systems (Cat. No. 98CH36187)*, Vol. 3, IEEE, 195–198
- Peng, H., & Bai, X. 2019, *The Journal of the Astronautical Sciences*, 1
- Powell, M. J. 1964, *The computer journal*, 7, 155
- Price-Whelan, A. M., Sipőcz, B., Günther, H., et al. 2018, *The Astronomical Journal*, 156, 123
- Reyes, J., & Cone, D. 2018, in *The Advanced Maui Optical and Space Surveillance Technologies Conference*
- Santurkar, S., Tsipras, D., Ilyas, A., & Madry, A. 2018, in *Advances in Neural Information Processing Systems*, 2483–2493
- Schildknecht, T. 1994, *Geod.-Geophys. Arb. Schweiz*, No. 49., 49
- Shell, J. R. 2010, *Optimizing orbital debris monitoring with optical telescopes*, Tech. rep., Air Force Space Innovation and Development Center Schriber AFB CO
- Shi, Y., & Eberhart, R. 1998, in *1998 IEEE international conference on evolutionary computation proceedings. IEEE world congress on computational intelligence (Cat. No. 98TH8360)*, IEEE, 69–73
- Shuster, M. D. 1993, *Navigation*, 8, 439
- Simonyan, K., & Zisserman, A. 2014, arXiv preprint arXiv:1409.1556
- Spurbeck, J., Jah, M., Kucharski, D., Bennett, J. C., & Webb, J. G. 2018, in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*
- Swietojanski, P., Ghoshal, A., & Renals, S. 2014, *IEEE Signal Processing Letters*, 21, 1120
- Wang, Z., Yan, W., & Oates, T. 2017, in *2017 international joint conference on neural networks (IJCNN)*, IEEE, 1578–1585
- Westlund, H. B., & Meyer, G. W. 2002, in *Graphics Interface*, Vol. 2002, 189–201
- Wetterer, C. J., & Jah, M. K. 2009, *Journal of Guidance, Control, and Dynamics*, 32, 1648
- Willison, A., & Bédard, D. 2016, *Advances in Space Research*, 58, 1318
- Zamek, S., & Yitzhaky, Y. 2006, *JOSA A*, 23, 3106
- Zhang, T., Xie, L., Li, Y., et al. 2018, *Energies*, 11, 3435
- Zhao, B., Lu, H., Chen, S., Liu, J., & Wu, D. 2017, *Journal of Systems Engineering and Electronics*, 28, 162