Space Object Identification Using Deep Neural Networks

Ian McQuaid
Air Force Institute of Technology
Ian.mcquaid.2@us.af.mil

Laurence D. Merkle
Air Force Institute of Technology
Laurence.Merkle@afit.edu

Brett Borghetti
Air Force Institute of Technology
Brett.Borghetti@afit.edu

Richard Cobb
Air Force Institute of Technology
Richard.Cobb@afit.edu

Justin Fletcher
Air Force Research Laboratory/Directed Energy Directorate
justin.fletcher.8@us.af.mil

ABSTRACT

The tracking of objects in orbit around the Earth is vital for both national security and the economic use of space. Much of the US space surveillance network's space object tracking activity comprises ground-based non-resolved optical observations of resident space objects in geosynchronous Earth orbit. To track orbiting objects, observations of the object must be separated in angle; in turn, observations must be separated temporally and associated. Associating observations with an object using only astrometric data (e.g. topocentric angular position) is complicated when an object maneuvers between observations or when multiple objects are in close proximity. To compensate for these complications, temporal variation in apparent object brightness across observations may be exploited. In this work, we formulate object-observation association as a pattern classification task and assess the suitability of convolution neural networks when applied to this classification task in simulation. Model performance is evaluated in three object-observation association scenarios: 1) no maneuvering, 2) objects maneuver in progressively closer proximity, and 3) objects permute stations with one another. In each simulation, photometric observations are generated as if observed by the Maui GEODSS site. This research contributes a foundational proof of concept for object classification via photometric signature.

1. INTRODUCTION AND MOTIVATION

Space situational awareness (SSA) is a mission area and problem domain comprising tasks which, if completed, maximize information about the state and intent of objects in orbit around the Earth. SSA is foundational to the protection of strategic national assets in space and peaceful economic use of space. The security of resident space objects (RSOs) in geosynchronous Earth orbit (GEO) is of elevated priority. Observation of GEO RSOs is particularly challenging, owning to the distance between the sensor and the RSO. Reaction to events is delayed, as transit to GEO takes longer than to lower altitude orbital regimes. GEO RSO observations are often collected using ground-based electro-optical sensors and consist of non-resolved images.

Tracking an RSO requires numerous observations; these observations must be separated in time to better approximate the orbit. In practice, RSOs must be observed at an interval of hours to guarantee sufficient angular

diversity for orbit estimation. Operationally, RSO tracking is complicated by sensor scheduling concerns such as target priority and exclusions, which can lead to substantial observation separation intervals. Temporal separation of observations introduces an additional complication: object-observation association. Objects may have maneuvered between observations. If an observation is associated with the orbital estimate build from observations of a different satellite, the resultant updated orbital estimate will be incorrect. Association is sometimes called correlation, and observations and tracks without a known “parent" RSO are sometimes called uncorrelated observations and tracks (UCOs and UCTs). In this work, we define association to be the problem of labeling an observation with either the object to which it corresponds, or a set of observations of the same object.

Current association techniques are effective in catalog maintenance use-cases. The literature contains a myriad of association methods [1], [2], [3], [4]. These techniques share a common limitation: all make exclusive use of astrometric information. That is, they make their decisions based on the observed kinematics of an RSO. For non-resolved GEO imagery means angular position in a topocentric coordinate system. Associating observations using only angles information has two main limitations: First, as RSOs decrease in apparent separation, uncertainty in orbital estimation eventually exceeds the metric accuracy of the observation, rendering differentiation statistically difficult. Second, there is no means to maintain accurate object identity differentiation through an apparent convergence. Analogously, if two objects were to switch apparent locations between observations, it would be impossible to determine this had occurred using only astrometric information.

There is, however, additional information available to observers of RSOs, which could be considered when performing object-observation association. One such source of information is the time-series of apparent brightness of an object, known as a photometric light curve, or photometry [5]. In this work, a photometry-based association technique is described, and the performance of this technique is reported. A deep neural network (DNN) is trained to extract relevant features from a photometric signature, and to map those features into classes corresponding to object identity.

2. BACKGROUND

A. Space Situational Awareness

SSA comprises many individual requirements, each having separate, often mutually-exclusive metrics to optimize. Coverage (the proportion of space near the Earth that is actively monitored), detection (the discovery of RSOs), tracking (the orbit estimation of RSOs), and characterization (understanding the function of RSOs) are a few such metrics. In this work, we consider a technique relevant to the GEO track creation process, the observation-track association task, and photometric analysis. An introduction to each is given in the following sections.

1) **GEO Track Creation**: Notionally, a track (also called an ephemeris) is an estimation of an RSOs position and velocity vectors. Before a track can be created, an RSO must be detected. Initial detection is often accomplished via optical sensors. These sensors report the light intensity they capture over each pixel in their field of view. One might then attempt to detect objects via radar instead of optical telescopes. However, radar systems that have frequency and power sufficient to detect a GEO RSO also have angle uncertainties sufficiently large as to limit the utility of those detections. Thus, the detection and tracking GEO RSOs is typically done using nonresolved imagery from electro-optical telescopes [6].

The manner in which this is done differs by sensor, though here GEODSS will be used as an exemplar [7]. A detection or observation (not to be confused with the machine learning term of the same name) is a single light streak usually several seconds long and consisting of several frames of data. Each frame consists of the pixel (i.e. CCD) outputs, where the units are usually count of electrons contained in each pixel. Longer exposure times enable the detection of dimmer objects. The dimmest object that a sensor can detect is referred to as the sensitivity of the sensor. Longer exposures also give the object more time to move around within a given pixel, as well as between pixels, resulting in more astrometric uncertainty. This can be mitigated by operating the sensor in rate-track mode.

A set of observations can be constructed into an object track via either performing an initial orbit determination from angle estimates, or by improving upon an existing ephemeris using the newly obtained angle measurements. In general, the principle is that greater angular separation in observations corresponds to more accurate ephemerides. However, nonresolved observations must first be associated to one another before a track can be created from them.

2) **Observation and Track Association**: Several observations of a given RSO must be aggregated before they can be used for track construction. This necessitates either collecting observations over long periods of continuous observation, or revisiting the RSO periodically. Continuous tracking is often impractical, so periodic revisiting is generally necessary. Periodic revisiting introduces the problem of determining which known tracks new observations correspond to, or which observations collectively correspond to the same untracked RSO. This problem is termed association, with the former being known as observation-to-track association and the latter observation-to-observation association. Association is a prerequisite to orbit determination and RSO tracking, and is thus foundational to SSA.

Association is a well-studied problem in the SSA community, with numerous approaches in active exploration. Most simply, a new observation may be associated to the cataloged RSO closest to that observation. Closest, is defined here to be the shortest distance between the RSO and the observation in the sensor's own azimuth and elevation topocentric coordinate system. This will yield correct associations in many situations. In the cases where it does not, however, more sophisticated techniques are required. The launch of new RSOs, or the creation of debris after a collision are a focus of the community.

3) **Photometry**: Optical observations contain brightness information, in addition to angular position. The intensity of light over time, also known as the photometric light curve or just photometry, is captured by the sensor, and may be exploited to reveal satellite state information. Kuroskai demonstrates spinning satellites that yield recognizable light curves, in that their spin-rates can clearly be determined with a transformation of the light curve signal to the frequency domain (i.e. FFT) [11]. Rodriquez specifically examines assumptions of object albedo, or material composition. They discovered that both composition and shape drive the form of the observed light curve, and attempted to construct a preliminary taxonomy of their photometric signatures. Ultimately the goal is to characterize the kinds of light curves one would expect from objects of a variety of compositions, shapes, and sizes [12].

It is possible to extract periodicity, average brightness, and a rough estimate of object size based on a light curve. Theoretically, the photometric light curve is affected by a number of factors, to include sun-object angle, object-sensor angle, and the object's incident illumination. Most importantly, the light curve is a function of the object's attitude and body shape. The process of estimating either of these is called light curve inversion, and has been studied extensively as it relates to determining the shape of asteroids [13] [14]. Hall specifically approaches the problem of characterizing attitude and body shape for an arbitrary object from its light curve, and comes to the conclusion that they should be extracted separately in order to obtain the best possible accuracy and tractability. He also notes that light curve inversion is based on the assumption that the RSO's geometry is convex. This is a generally accepted assumption when the process is applied to asteroids, though it may not be appropriate for man-made satellites. That being said, the impact of the object convexity assumption is unknown, though it limits the current practical application of light curve inversion [15].

B. Deep Neural Networks

Machine learning is a set of techniques that automatically develop algorithms to map inputs (called observations) to outputs (called labels), given data. When applying deep learning to a problem, some relationship between inputs and outputs is assumed. The form and quantification of that relationship is, however, unknown. Machine learning techniques extract that relationship automatically, in pursuit of the minimization of some quantified loss function. The algorithms, hereafter referred to as models, are the product of a search through the space of potential relationships.

1) **Classification**: Supervised, feed-forward deep learning models broadly fit into two categories based on the kind of label they predict. Classifiers predict categorical labels, and regressors predict continuous numeric labels. RSO identity is a categorical variable, and is thus compatible with classification. The simplest case is the two-class problem in which the model provides a Boolean response to each observation. Multi-class problems can be viewed as simple extensions of the two class problem by viewing an observation's label as a "one hot" vector. Instead of making the label a natural number I indicating that an observation belongs to class i, the label is instead a Boolean vector (e.g. (0, 0,…, 0,1,0,…,0)) containing a single 1 in the ith position and n-10's in the remaining positions, where n is the number of classes. This is equivalent to simultaneously answering n two-class problems, where problem i is to determine if an observation is either class i, or any class j where $j \neq i$ [16].

2) **Artificial Neural Networks**: Artificial neural networks are a subset of machine learning models often referenced in modern (or deep) machine learning. The fundamental element of an ANN is a neuron, motivated by the biological cell of the same name. The inputs to the virtual neuron are each individually multiplied by “synapse" weights and the results summed. This value is given as input to a non-linear activation function that produces the neuron’s output activation. The synapse weights are the trainable parameters within the model, and the non-linear activation function permits compositions of neurons to produce complex behavior, rather than behavior that could be obtained from a single neuron. Theoretically, provided one uses non-linear activation functions, ANNs could replicate the mapping of any training set, provided the network contained a sufficient number of neurons [17].

At the time of this writing, the state-of-the-art approach to ANN design is to limit the width of each intermediate layer, but to include many layers (hence “deep") [17]. Such networks generally have larger capacity (number of trainable weights) than needed, and thus tend to overfit. Several regularization techniques exist to combat this. Batch normalization serves to counteract the impact of updating layers sequentially. Dropout is another popular regularization technique that randomly removes a subset of neurons during each training batch. This is achieved by examining the output activation of each neuron and stochastically, at some preset frequency, fixing those neuron’s activations to 0 [17]. This restriction works to counteract specific local patterns that form during fine-tuning.

3) **Convolutional Neural Networks**: At present, computer vision is dominated by deep learning. Convolutional neural networks (CNNs), specifically, have become the de facto standard for most image processing tasks. The popularity of CNNs owes to the automation of feature extraction they provide. CNNs learn a parameterized mapping from pixels to relevant features automatically via gradient descent. It is impractical to parameterize each pixel-to-feature mapping of an image, so CNNs implement weight sharing. This is realized via the parameterization of convolutional kernels, which are then convolved with the pixels of an image to produce a feature map. A CNN applies several independently-trainable instances of these transformations to the image, producing several channels of features. The weights parameterizing the convolutional kernals are trained via gradient descent to minimize a cost function. These convolution layers are typically stacked to increase the capacity of the model, and combined with pooling operations to keep the model size manageable.

C. *Related Work*

Linares [20] trained DNNs to classify observations of RSOs into one of four categories: rocket bodies, controlled payloads, uncontrolled payloads, and debris. In this work, all RSOs were GEO objects with identical orbits. The DNNs trained in that study achieved a classification accuracy of 99:6%, indicating the potential utility of DNNs in categorizing RSOs based on observed light curves. This work may be seen as an extension of that approach; in this work RSOs are simulated in different orbits and the DNN is trained to map from observation to object identity, rather than object type.

Jia [21] examined several different machine learning techniques (namely k nearest neighbors, convolutional neural networks, and recurrent neural networks) at space object identification from photometric observations. Typically this requires simulation in order to produce the quantity of observations needed to train the machine learning models, however this introduces questions as to the accuracy of that simulation. Instead Jia chose to use real world data for different stars, namely data from the Optical Gravitational Lensing Experiment (OGLE). They then noted good performance (better than 94% accuracy) across all methods, though with the DNN methods performance was better than 96%. This lends additional evidence towards the effectiveness of DNNs in mapping photometry to object identities, though it remains to be seen how effective this technique is when the quantity of data available is reduced.

3. APPROACH AND METHODOLOGY

There is, however The objective of this study is to demonstrate the feasibility of association GEO RSO observations using deep neural networks. To that end, this work hypothesizes that deep neural networks (DNNs) trained on photometry can be applied to GEO SSA to address current UCO and UCT association ambiguities. We consider three lines of inquiry to evaluate this hypothesis:

- Can a DNN be trained to map photometric observations to the corresponding RSO?
- Can a DNN be trained to map photometric observations to the corresponding RSO when the RSOs are closely spaced?

- Can a DNN be trained to map photometric observations to the corresponding RSO when the RSOs are maneuvering?

An experiment is conducted to assess each of these questions. All three experiments share common elements: the dataset generation, the DNN topology, and the HPC resources leveraged. These are described first, followed by individual descriptions of the experiments.

A. *Dataset Generation*
To create a simulated dataset representative of real-world phenomenology, several physical phenomenon must be modeled. AGI's STK [22] will be used to model that phenomenology. Satellites are selected from those in the Maui GEODSS site's field of view (FOV), such that RSOs at relatively different stations. A catalog of GEO RSOs taken from Space-Track.org on 31 July 2017 [23], is filtered to meet these constraints, yielding a subset of 8 RSOs. These RSOs are JCSAT 16, INTELSAT 907, GOES 13, AMC-16, GOES 16, DIRECTV 11, MEXSAT 3, and GALAXY 15. The scenario used started on 1 Jan 2019 at 1200 UTC and lasted until 1 Jan 2020 at 1200 UTC.

STK allows the specification of an RSO's photometric properties such as material, attitude profile, and geometry. STK models the entire surface of the RSO as a single material. This models the spectral response, meaning the wavelengths of light absorbed or reflected, for that RSO. The spectral response curves are internal to STK and are based on NASA Reference Publication 1121 [24]. The attitude profiles considered in this research are nadir alignment with orbit normal constraint, and nadir alignment with Sun constraint. The former aligns the RSO z axis to the nadir direction (opposite the geocentric position vector) and constrains the RSO's x axis to be in the direction of the orbit normal vector. The second profile constrains the RSO's z axis to point in the nadir direction and the RSO's x axis to point towards the Sun. Finally, the shapes used are the GEOComm and LEOImaging shapes, which are internal to STK.

It is also necessary to model a sensor to generate the dataset observations. Photometry is highly sensitive to the observer's CCD, so the choice of sensor is important. Here the GEODSS telescopes and CCD(s) were used. The operational GEODSS sensors are located at three sites: Maui, HI (20.71° N by 156.26° W), Diego Garcia (7.41° S by 72.45° E), and Socorro, NM (32.82° N by 106.66° W). This study focuses on the Maui GEODSS site specifically. The photometric observations are generated from STK using the EOIR plug-in taking into account the object properties and attitude as well as the GEODSS sensor. The generated photometry only consists of a single channel (from 0:4-0:9µm).

For simulation purposes, scheduling concerns are ignored and it is instead assumed that the sensor can observe everything in its field of view simultaneously. A 6 hour minimum time between observations is imposed so that no object is observed unrealistically often. However, the sensor is not free to observe an object at any time. In order to create an observation, the sensor must be in darkness (penumbra or umbra) and the target in sunlight (direct sun). This means that, in practice, the simulation will exhibit revisit rates lower than 1 observation every 6 hours.

The scenario begins nearly 2 years after these space-track ephemeris are dated, and lasts an additional year. As such, SGP4 will not yield a representative ephemeris. This is because operational GSO satellites must keep station to remain in compliance with international law. SGP4 will not model this by default, nor should it. Station keeping burns are necessary, and must be modeled for our observations to be representative.

B. *Deep Neural Network Topology*
The machine learning task under consideration is a supervised two-class pattern classification problem. Models will be trained on photometric observations of two RSOs, where the classes the model must predict correspond to the identity of those two RSOs. Trained models map a given photometric observation to the RSO to which the observation corresponds. The loss function used is the binary cross-entropy function given by

$$-\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)],$$

where N is the number of observations, yi the true class label, and yi hat the class label predicated by the model. The class label will be either 0 or 1, which is arbitrarily assigned to the two RSOs under consideration [25]. The main performance metric monitored is accuracy, meaning the percentage of observations mapped to the correct RSO.

To accomplish this task, an appropriate network topology is needed. The topology used here is built with five successive modules of two 1D convolution layers with kernel length 3 and 64 filters followed by a max pooling layer with pool size 2. Afterwards the outputs are flattened and given as input to 3 fully connected layers of width 64 before the output layer. All non-output layers use the rectified linear unit (ReLU) activation function, whereas the output layer uses the sigmoid activation function. A graphical depiction of this topology is shown in Figure 1.

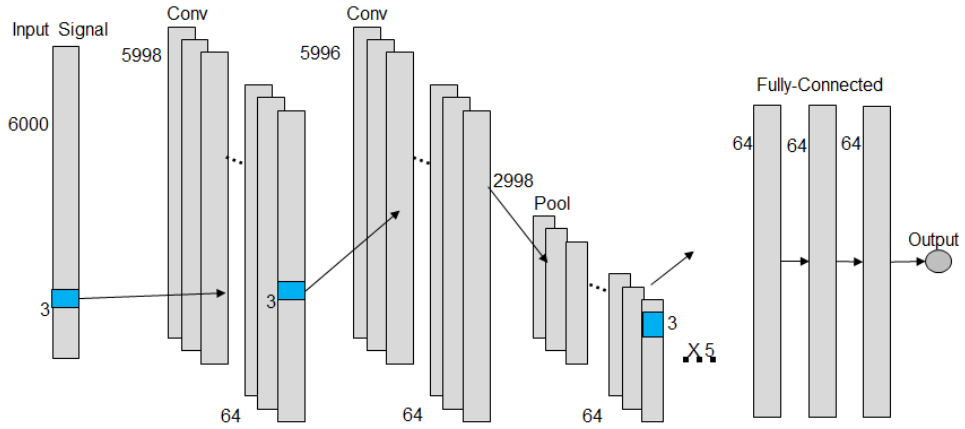

Fig. 1. Diagram of convolutional neural network topology used to associate photometric observations with RSO identity

Network training will continue until 30 epochs pass without improvement to validation performance. The intent is to account for any variance in the amount of training needed autonomously so that many models can be executed simultaneously and without human tuning. The batch size was set to 128 observations per batch in an attempt to find a trade-off between runtime and randomness of training batch contents. The Adam optimizer was used and default learning rate used with it. This ANN was implemented, trained, and tested using the Keras Python library [25].

C. *Baseline Association of RSOs*

This experiment addresses the first research question: Are there any cases where a DNN can use photometry to correctly associate observations to the corresponding RSO? The machine learning task used to answer that question is pairwise classification of 8 GEO RSOs where those RSOs are all station keeping and spaced apart from one another. These 8 RSOs will have different geometries, materials, and attitude profiles. Table I contains the photometric properties of each of the 8 RSOs in this scenario. This problem is perfectly and easily solved via metric data, as even the nearest neighbor method ought to produce correct association in such a benign case. However it is doubtful that photometric observations will be easier to associate in other cases, so this serves as a sort of theoretical maximum performance. A geographic depiction of the catalog used for this experiment is shown in Figure 2.

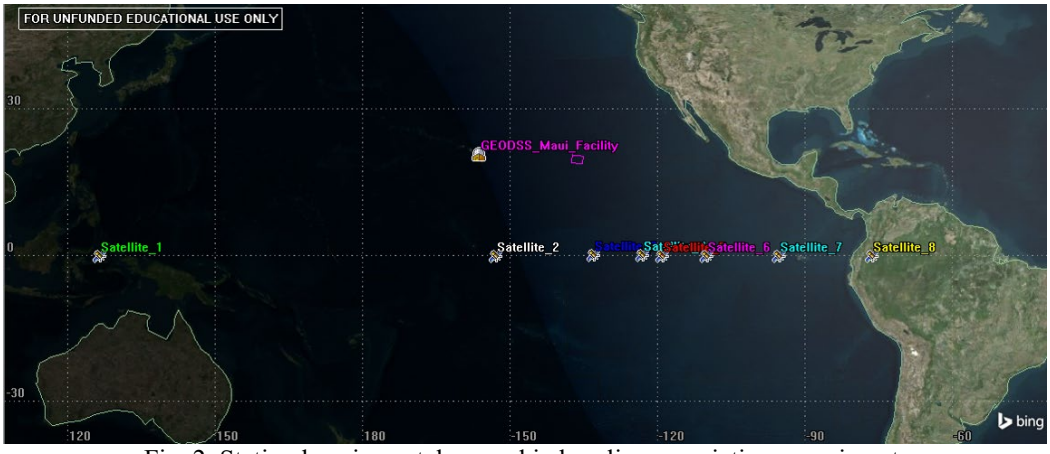

Fig. 2. Station keeping catalog used in baseline association experiment.

TABLE I RSO SHAPES, MATERIALS, AND ATTITUDE PROFILES USED IN THIS EXPERIMENT.

RSO Name	Original Name	Attitude Profile	Material	Geometry
Satellite 1	JCSAT 16	1	Aluminum MLI	GEOComm
Satellite 2	INTELSAT 907	1	Aluminum MLI	LEOImaging
Satellite 3	GOES 13	1	Solar Panel	GEOComm
Satellite 4	AMC 16	1	Solar Panel	LEOImaging
Satellite 5	GOES 16	2	Aluminum MLI	GEOComm
Satellite 6	DIRECTV 11	2	Aluminum MLI	LEOImaging
Satellite 7	MEXSAT 3	2	Solar Panel	GEOComm
Satellite 8	GALAXY 15	2	Solar Panel	LEOImaging

Attitude Profile 1 is “Nadir align w/ orbit normal const."
Attitude Profile 2 is “Nadir align w/ sun const."

These 8 RSOs are considered pairwise so that the ANN is only ever trained on a two class problem. Note that there are 28 RSO pairings considered here. The intent is to determine which pairs of RSOs are most difficult to distinguish in an attempt to understand the utility and capability of photometric light curves.

For a specific RSO pair of the 28 pairs under consideration, 3-fold cross validation is used with 25% of the training fold observations used for the validation set. This process will be replicated 10 times with different random seeds so that 30 performance estimates are available.

With a dataset generated and experiment executed, the results will be aggregated and presented. The intent is to determine in which cases one can reasonably expect RSOs to be distinguishable, at least provided the quantity and quality of data generated here. Armed with this knowledge, this research will then assess the utility of photometric data in distinguishing RSOs in close proximity to one another.

D. *Association of RSOs with Varying Proximity*
This experiment addresses the second research question: Can a DNN use photometry to associate observations of RSOs even if those RSOs are closely spaced? The machine learning task used to answer that question is classification of one specific GEO RSO against each of 7 others. The intent is to determine if proximity has an impact on the photometric association process. The RSOs in this scenario will all station keep at increasing distances; the stations assigned to each are shown in Table II. A graphic showing this catalog can be found in Figure 3.

TABLE II LONGITUDE OF ASSIGNED STATION FOR EACH RSO IN THE VARYING RSO PROXIMITY SCENARIO.

Satellite #	1	2	3	4
Longitude	−150.0°	−150.01°	−150.1°	−150.2°

Satellite #	5	6	7	8
Longitude	−150.3°	−150.4°	−150.5°	−150.6°

Satellite 1 will serve as the baseline objects in this experiment and it will be examined against satellites 2 through 8. This amounts to 7 different two-class classification problems: Satellite 1 vs Satellite i for i = 2,3,…,8. The photometric qualities of these RSOs will be selected so that satellite 1 should be distinguishable from the remainder of the catalog. The satellites other than satellite 1 will all share the same photometric qualities so that the only factor

varying in this experiment is proximity. As for specific values used for those photometric qualities, they will be informed by the results of the baseline experiment.

For each of these individual two-class problems, 6-fold cross validation will be employed and will be replicated 5 times so that there are 30 performance estimates for each problem. These results will then be collected and analyzed.

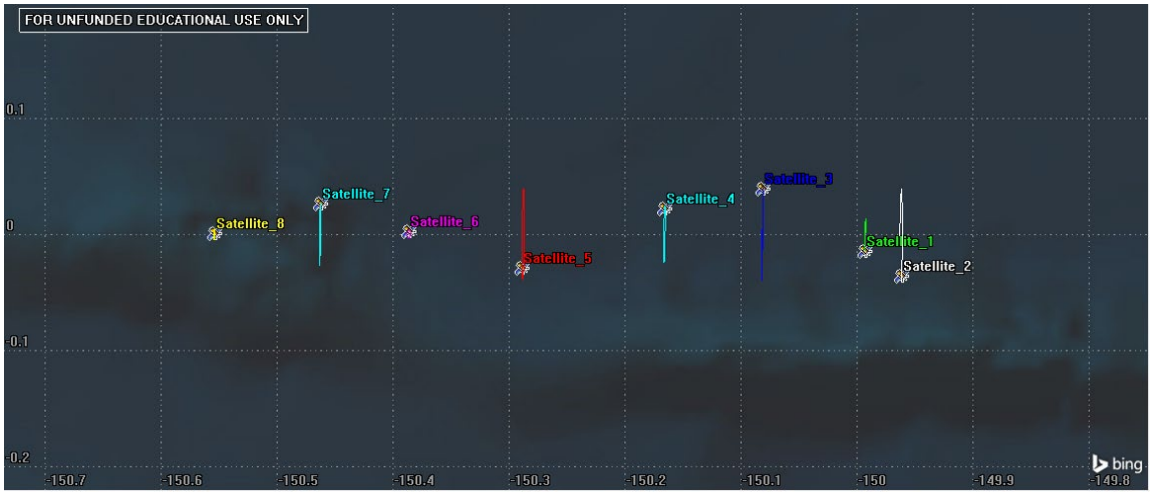

Fig. 3. Station keeping catalog used in varying RSO proximity experiment.

With the results of this experiment in mind, one should be able to determine if RSO proximity has a significant impact on the effectiveness of this association technique. It is reasonable to anticipate that the photometric signature would remain just as salient a feature in object differentiation regardless of proximity. With that in mind, one would expect a high, though likely not 100%, accuracy between each assessed RSO pair, and that that accuracy does not degrade with proximity. If RSOs could be differentiated in the previous experiment that trend is anticipated to continue here.

E. *Association of Maneuvering RSOs*

This experiment addresses the third research question: Can a DNN use photometry to associate observations of RSOs even if those RSOs are maneuvering? The machine learning task used to answer that question is pairwise classification of 8 GEO RSOs where those RSOs will periodically permute stations with one another. The generation of the dataset itself will closely follow that used in the baseline described earlier in this chapter. The scenario will be a year long with observations taken every 6 hours (as constrained by lighting of the sensor and RSOs) and the same shapes, materials, and attitude profiles will be used. The difference is that Astrogator is not used to command the RSOs to keep station throughout the whole year. Instead every 31 days the RSOs will change to a station occupied by a different RSO. The schedule is given by the ordered list: JCSAT 16, INTELSAT 907, GOES 13, AMC-16, GOES 16, DIRECTV 11, MEXSAT 3, and GALAXY 15. Each time a given RSO maneuvers it will advance one position in a randomized version of the above list.

The way that Astrogator accomplishes this is not dissimilar to the way it kept station in the baseline scenario. The target sequence used will still amount to a Hohmann transfer, though the tolerance is initially loosened to allow Astrogator to converge for large-scale maneuvers. The reason being is that, due to the scale of these maneuvers, converging on appropriate maneuvers can be difficult. Progressively letting Astrogator step closer to the desired station ensures that tight tolerances are only expected when maneuver magnitude is also relatively small. This is done in four successive target sequences, with tolerances equal to 10, 5, 1, and 0:05 longitude, respectively. A graphic showing the resulting catalog can be seen in Figure 4.

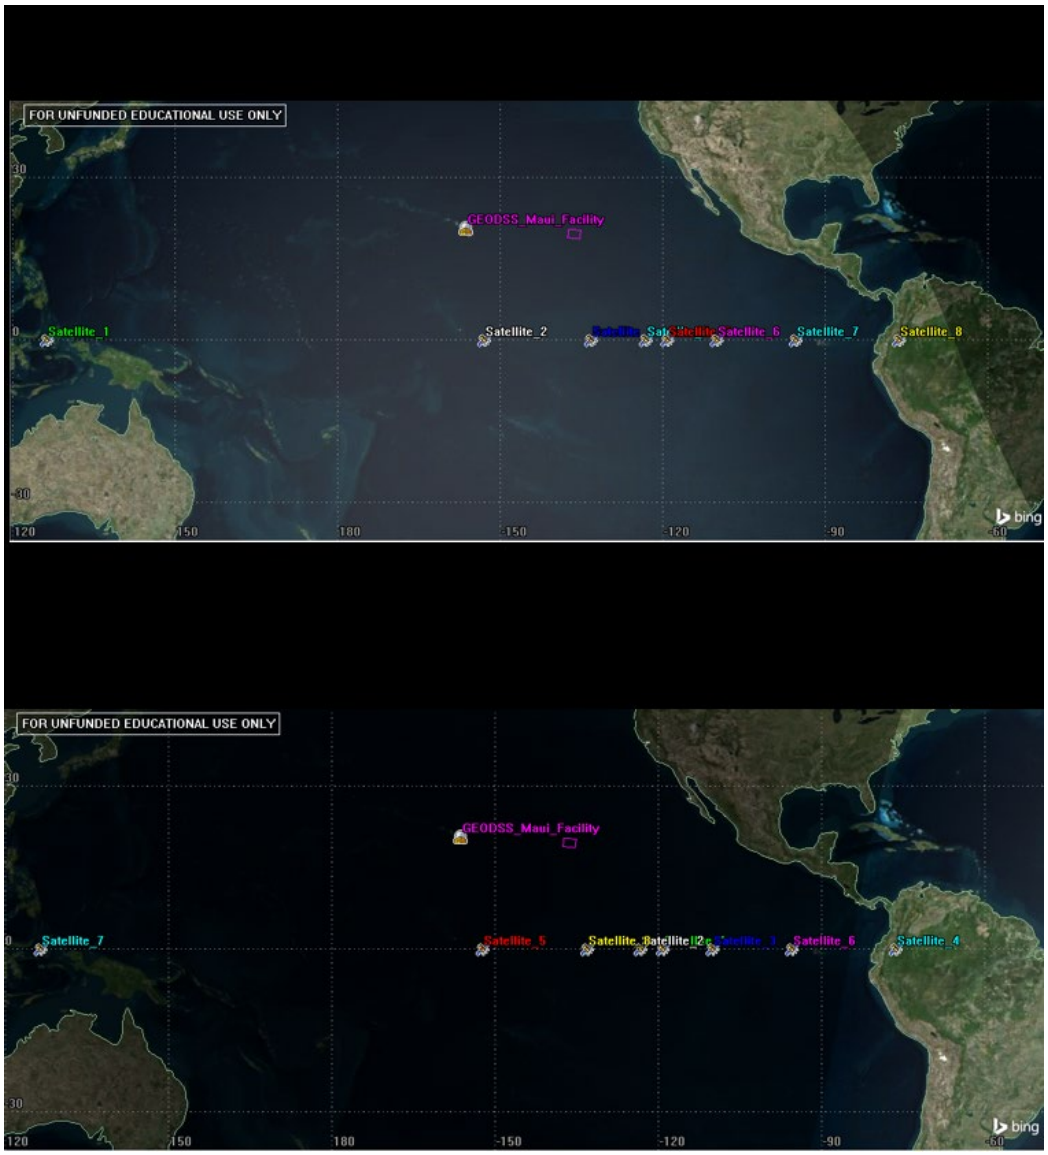

Fig. 4. Maneuvering catalog where RSOs permute stations every 31 days. Used in the maneuvering experiment.

Similar to the baseline experiment, this experiment considers every pair of RSOs (28 in total). For each pair, 3-fold cross validation is used with 25% of the training folds used for validation. Then 10 replications of the training and evaluation are run to yield 30 performance estimates. As in the previous experiment, one would not expect the introduction of maneuvers to make differentiating RSOs an easier task. RSOs that were not easily distinguished in the station keeping experiment are unlikely to be easily distinguished in this experiment. For those RSOs that could be accurately distinguished, it is suspected that maneuvering and the resulting change in attitude will make this a more difficult task than in the previous experiments.

4. RESULTS

The experiments described in the previous section were executed and the results collected. The following sections contain a description of these results, as well as a brief discussion of their implications and outstanding or unexpected findings. As previously, these sections correspond to a baseline station keeping scenario, a varying RSO proximity scenario, and finally a maneuvering scenario.

A. *Baseline Association of RSOs*

Regarding the station keeping scenario, recall that this experiment consisted of a catalog of 8 station keeping GSO RSOs. One year of simulated photometric observations were generated for those RSOs, and a deep convolutional neural network was evaluated on each RSO pairing. The primary metric used in evaluation of these ANN classifiers is accuracy, that is, the percentage of correct labels the classifiers produced on the respective testing folds. Recall that each RSO pair will have 30 such accuracy values. The mean and standard deviation of those values are presented in Table III. Each cell represents the mean or standard deviation for the classifiers that compared satellite i to satellite j where i and j are the row and columns numbers of that cell, respectively. This is why the diagonal is excluded (no RSO is ever compared to itself). Similarly, the values below the diagonal are aggregated with those above the diagonal and are thus excluded (the satellite 2 vs satellite 1 cell is combined with the satellite 1 vs satellite 2 cell).

TABLE III ACCURACY MEANS AND STANDARD DEVIATIONS FOR EVERY RSO PAIRING IN THE BASELINE ASSOCIATION EXPERIMENT. EACH CELL SUMMARIZES A SAMPLE OF SIZE 30.

Mean Accuracy

	Satellite 1	Satellite 2	Satellite 3	Satellite 4	Satellite 5	Satellite 6	Satellite 7	Satellite 8
Satellite 1	-	67.19%	82.25%	89.45%	87.77%	89.60%	76.94%	81.64%
Satellite 2	-	-	100.00%	99.25%	99.63%	93.95%	97.96%	64.85%
Satellite 3	-	-	-	100.00%	54.76%	98.61%	52.76%	99.94%
Satellite 4	-	-	-	-	99.66%	54.59%	97.24%	70.96%
Satellite 5	-	-	-	-	-	95.85%	62.86%	98.17%
Satellite 6	-	-	-	-	-	-	95.81%	77.99%
Satellite 7	-	-	-	-	-	-	-	99.57%
Satellite 8	-	-	-	-	-	-	-	-

Standard Deviation of Accuracy

	Satellite 1	Satellite 2	Satellite 3	Satellite 4	Satellite 5	Satellite 6	Satellite 7	Satellite 8
Satellite 1	-	10.85%	7.63%	11.87%	1.92%	8.42%	2.57%	6.17%
Satellite 2	-	-	0.00%	3.83%	0.51%	11.13%	1.44%	8.54%
Satellite 3	-	-	-	0.00%	4.00%	7.46%	4.69%	0.08%
Satellite 4	-	-	-	-	0.21%	1.71%	10.16%	10.60%
Satellite 5	-	-	-	-	-	8.62%	9.00%	0.47%
Satellite 6	-	-	-	-	-	-	10.39%	1.96%
Satellite 7	-	-	-	-	-	-	-	0.87%
Satellite 8	-	-	-	-	-	-	-	-

From these results, a somewhat non-intuitive pattern appears to emerge. Recall that 50% accuracy is the worst one should expect in a 2-class problem, so the cases in Table III with means near 50% correspond to classifiers failing to significantly tell the RSOs in question apart. As far as a proof of concept, certainly some RSO pairings appear quite distinguishable (e.g. Satellite 2 vs Satellite 3). That is to say that it would appear there are cases where photometric observations are a reliable way of distinguishing RSOs. It should be noted that any model trained on observations of satellite 1 performs poorly, relatively speaking. The reason is not immediately obvious, though if one recalls Figure 2, Satellite 1's station is fairly distant from that of the rest of the catalog. This could lead to significantly different viewing aspects, or exposed facets, amounting to a photometric signature indistinguishable from the rest of the catalog. With all that in mind, what causes some cases to work better than others is hopefully made more clear in Figure 5.

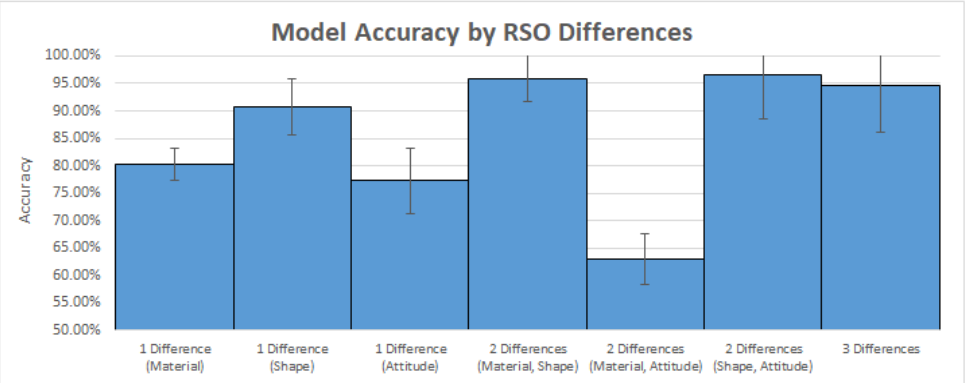

Fig. 5. Mean accuracy arranged by manner in which each RSO pair differs. Bar height gives mean accuracy given those differences, and error bars show standard deviation. Note that each bar describes 4 RSO pairs, hence a sample of size 120.

Figure 5 aggregates the samples summarized in Table III based on the way the RSOs differ. So, for instance, every RSO pair that differ only in material are combined in the first bar. This would be satellite 1 vs satellite 3, satellite 2 vs satellite 4, satellite 5 vs satellite 7, and satellite 6 vs satellite 8 (see Table I). Each bar will have 4 RSO pairs included. Since each RSO pair is itself a sample of size 30, each bar describes a sample of size 120. The height of the bar is the mean accuracy of that sample, and the error bars correspond to the sample standard deviation.

With that description in mind, the generally decent performance of this technique can be seen. Aside from the 2 differences - material and attitude case, most bars seem accurate. Of the cases that don't perform as well, they all hold RSO shape constant. That is, the performance of this technique appears to degrade in cases where RSOs have the same geometry, even if the material and/or attitude profile varies. This would agree with intuition. Even though it was anticipated that all of these factors would impact an RSO's photometric signature, this is evidence that geometry has the most significant impact.

B. *Association of RSOs with Varying Proximity*
When considering the subsequent experiment regarding RSO proximity, we start with RSOs that seem easily distinguished. In this case RSOs that differ in both material and shape are taken, as this case appears to perform very well in the station keeping experiment.

The specific RSOs used in the varying RSO proximity scenario are satellites 2 and 3. These differ in both geometry and material though have the same attitude profile. The results obtained can be seen in Table IV and are summarized in Figure 6. Since every RSO pairing wasn't considered in this experiment, 6-fold cross validation was used instead of 3-fold so that the resulting accuracy samples were still each of size 30.

TABLE IV ACCURACY MEANS AND STANDARD DEVIATIONS FOR EVERY RSO PAIRING IN THE VARYING RSO PROXIMITY EXPERIMENT. EACH CELL SUMMARIZES A SAMPLE OF SIZE 30.

Mean Accuracy								
	Satellite 1	**Satellite 2**	**Satellite 3**	**Satellite 4**	**Satellite 5**	**Satellite 6**	**Satellite 7**	**Satellite 8**
Satellite 1	-	87.04%	88.43%	88.93%	88.04%	87.90%	88.65%	87.76%
Standard Deviation of Accuracy								
	Satellite 1	**Satellite 2**	**Satellite 3**	**Satellite 4**	**Satellite 5**	**Satellite 6**	**Satellite 7**	**Satellite 8**
Satellite 1	-	3.10%	3.49%	2.77%	2.47%	3.26%	2.63%	3.15%

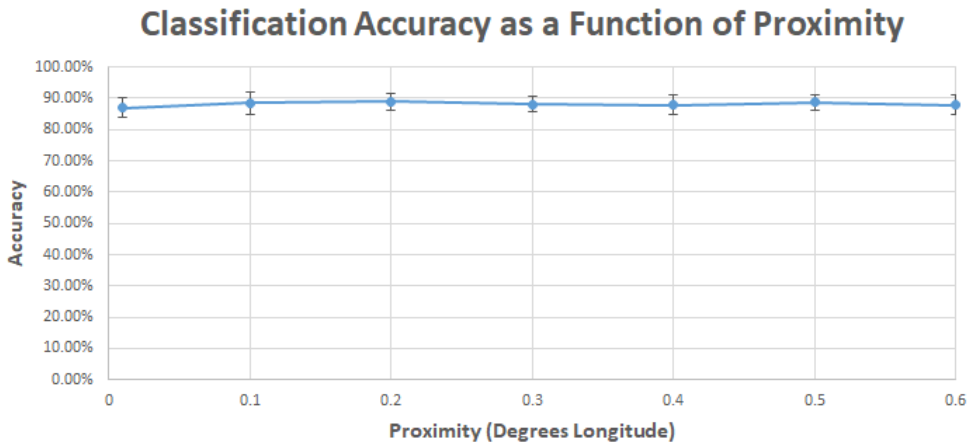

Fig. 6. Line graph showing change in the mean and standard deviation of accuracy as RSO proximity changes. Each data point summarizes a sample of size 30.

One notices that these performance estimates are all near that of the previous experiment. This is as predicted, as proximity should not impact the performance of this technique. To further strengthen that claim, there does not seem to be a pattern of improvement or degradation with RSO proximity. So even gradually increasing or decreasing RSO distance doesn't appear to impact this technique. At the extreme satellite 1 and 2 are nearly station sharing. All of that to say that, provided the RSOs were distinguishable initially, they still can be distinguished when placed close together.

C. *Association of Maneuvering RSOs*

For the maneuvering experiment the same reporting mechanisms are used as in the baseline experiment. The means and standard deviations of the resulting accuracy samples can be found in Table V. The table is interpreted in the same manner as table III.

TABLE V ACCURACY MEANS AND STANDARD DEVIATIONS FOR EVERY RSO PAIRING IN THE MANEUVERING ASSOCIATION EXPERIMENT. EACH CELL SUMMARIZES A SAMPLE OF SIZE 30.

Mean Accuracy								
	Satellite 1	**Satellite 2**	**Satellite 3**	**Satellite 4**	**Satellite 5**	**Satellite 6**	**Satellite 7**	**Satellite 8**
Satellite 1	-	69.85%	52.58%	67.46%	63.42%	73.96%	60.26%	71.55%
Satellite 2	-	-	69.02%	50.67%	58.83%	50.83%	59.77%	57.72%
Satellite 3	-	-	-	67.45%	58.60%	72.42%	55.76%	72.93%
Satellite 4	-	-	-	-	58.72%	51.85%	58.64%	58.70%
Satellite 5	-	-	-	-	-	59.38%	50.19%	68.74%
Satellite 6	-	-	-	-	-	-	62.86%	53.91%
Satellite 7	-	-	-	-	-	-	-	67.22%
Satellite 8	-	-	-	-	-	-	-	-

Standard Deviation of Accuracy								
	Satellite 1	**Satellite 2**	**Satellite 3**	**Satellite 4**	**Satellite 5**	**Satellite 6**	**Satellite 7**	**Satellite 8**
Satellite 1	-	2.79%	2.04%	2.65%	3.17%	3.04%	2.87%	2.77%
Satellite 2	-	-	4.06%	1.04%	4.07%	1.11%	4.62%	3.48%
Satellite 3	-	-	-	5.26%	4.67%	4.75%	3.22%	3.73%
Satellite 4	-	-	-	-	7.08%	2.35%	4.22%	5.87%
Satellite 5	-	-	-	-	-	3.64%	1.09%	5.64%
Satellite 6	-	-	-	-	-	-	3.36%	3.64%
Satellite 7	-	-	-	-	-	-	-	3.83%
Satellite 8	-	-	-	-	-	-	-	-

At a glance, one notices the performance degradation due to the RSO maneuvers. Though some pairings remain somewhat distinguishable, many pairing yielded accuracies near 50%. Again, this means that those RSO pairs could not be distinguished any more accurately than a random guess. That maneuvering RSOs are more difficult to distinguish is not surprising, though perhaps the degree to which those maneuvers impact an RSOs photometric signature is. Figure 7 was constructed for the same purpose as in the station keeping experiment: to determine in what cases specifically DNNs trained on photometric-based observation association will be effective. Unfortunately, the differences between bars are mostly within the expected variances of each case. As far as viability, some cases perform significantly better than a 50% random classifier. Given these bars correspond to a sample of size 120, it is reasonable to conjecture that this is evidence of a learned relationship between the photometric signature and RSO identity. It is possible that more observations are simply needed from different perspectives in order to prepare the classifier for the dynamic behavior presented by maneuvering RSOs. If that were true, then one would expect performance to improve significantly as more training data was used.

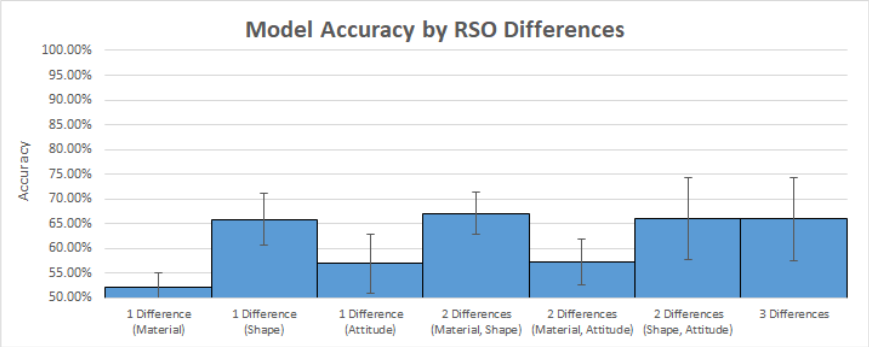

Fig. 7. Mean accuracy by RSO pairing differences. Bar height gives mean accuracy given those differences, and error bars show standard deviation. Note that each bar describes 4 RSO pairs, hence a sample of size 120.

In total, these results demonstrate the viability of DNNs trained on photometric observations to distinguish RSOs, in some cases, specifically that of RSOs with fundamentally different geometries. Furthermore this appears to be agnostic of RSO spacing so that, in those cases with different RSO geometries, this technique could provide the community with a valuable capability, as association of closely spaced objects is a known limitation of traditional astrometric-based techniques. Finally, for maneuvering RSOs the utility of DNNs trained on photometry appears significantly degraded. That being said, the technique yielded sufficient performance to warrant further research. It is possible that further training in more diverse scenarios could produce a classifier with similar performance to that seen in the station keeping scenario.

5. CONCLUSIONS AND FUTURE WORK

There is, however Results in this work indicate that there are cases where DNNs can use photometry to distinguish GEO RSOs. RSOs with different geometries that simply keep station are most likely to be distinguishable. Additionally, we find that DNNs can distinguish RSOs via photometry even when those RSOs are in relatively close proximity to one another. It would appear that, up to 0.01° at least, effectiveness of the DNN in this task does not degrade with proximity. Finally, in cases where the RSOs maneuver to permute positions with one another, the observed performance indicates that either the RSO's photometric signatures are insufficiently varied to enable distinction, or that insufficient data was captured to generalize to the contexts used in evaluation. Further investigation is needed to determine the root cause of the observed performance.

The assumptions and limitations of this research should be considered before generalizing any of the conclusions drawn in this work. As this data was simulated, the extension of the presented technique to real world data is not necessarily supported. Photometric signatures, in particular, are known to be very sensitive to geometry, atmosphere, etc., so that matching theory with observed values is not always straight forward. Thus, the procedure followed in this work should be evaluated on real data to determine if those conclusions generalize outside the simulation.

There are several avenues for continuation of this research. The generalization of an RSOs photometric signature is the most immediate such avenue. First, the ability to generalize in time, even in a non-maneuvering scenario, needs to be demonstrated before this technique is viable in any operational context. Additionally, generalization of the photometric signature to different maneuvers is a more difficult, though very interesting problem. It may be that photometric signatures, as framed here, are simply insufficient to accomplish the tasks described. A natural follow-on would be to consider multi-channel signatures. The most basic form of this would be to divide the visible spectrum into red, blue, and green light components. More extreme still would be to create a large number of channels with a very fine spectral resolution, potentially even expanding beyond the visible spectrum. Finally this technique should be applied to real world data. The phenomenology required to model photometric light curves is complex and sensitive. Ultimately this technique is only useful if it works on real world light curves, thus that is a natural extension of this research.

ACKNOWLEDGMENT

The authors would like to thank AFRL/RD for sponsoring this research, as well as MHPCC for access to their computing cluster in order to conduct the computational experiments described in this document. The views and conclusions contained in this document are those of the authors and should not be interpreted as rep- resenting the official policies, either expressed or implied, of the United States Air Force or the U.S. Government.

REFERENCES

[1] N. Singh, J. Horwood, J. Aristoff, A. Poore, and C. Sheaff "Multiple hypothesis tracking (MHT) for space surveillance: Results and simulation studies," 2013. [Online]. Available: http://www.dtic.mil/docs/citations/ADA591392

[2] J. Stauch, M. Jah, J. Baldwin, T. K. A. A, "Mutual application of joint probabilistic data association, filtering, and smoothing techniques for robust multiple space object tracking," arc.aiaa.org. [Online]. Available: https://arc.aiaa.org/doi/pdf/10.2514/6. 2014-4365

[3] I. Hussein, C. Roscoe, M. W. P, "Track-to-Track Association Using Bhattacharyya Divergence," amostech.com. [Online]. Available: https://www.amostech.com/TechnicalPapers/2015/Astrodynamics/Hussein.pdf

[4] K. Fujimoto, D. J. Scheeres, J. Herzog, and T. Schildknecht, "Association of optical tracklets from a geosynchronous belt survey via the direct Bayesian admissible region approach," Advances in Space Research, vol. 53, no. 2, pp. 295–308, jan 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0273117713007059

[5] M. Bass, C. DeCusatis, J. Enoch, V. Lakshminarayanan, G. Li, C. Macdonald, V. Mahajan, and E. Van Stryland, Handbook of Optics, Third Edition Volume II: Design, Fabrication and Testing, Sources and Detectors, Radiometry and Photometry, 3rd ed. New York, NY, USA: McGraw-Hill, Inc., 2010.

[6] D. C. Mark Ackermann, Rex Kiziah, Peter Zimmer, John McGraw, "A SYSTEMATIC EXAMINATION OF GROUND-BASED AND SPACEBASED APPROACHES TO OPTICAL DETECTION AND TRACKING OF SATELLITES," 31st Space Symposium, 2015.

[7] R. Weber, "The ground-based electro-optical detection of deep-space satellites," R. E. Franseen and D. K. Schroder, Eds., vol. 0143. International Society for Optics and Photonics, sep 1978, pp. 59–69. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding. aspx?articleid=1227351

[8] J. Siminski, O. Montenbruck, H. Fiedler, and T. Schildknecht, "Shortarc tracklet association for geostationary objects," Advances in Space Research, vol. 53, no. 8, pp. 1184–1194, apr 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0273117714000520

[9] J. Siminski, O. Montenbruck, H. F. A., "Best Hypothesis Search on Iso-Energy-Grid for Initial Orbit Determination and Track Association," core.ac.uk. [Online]. Available: https://core.ac.uk/download/pdf/18591938.pdf

[10] J. Siminski, H. Fiedler, T. S. P. of the 6th, and undefined 2013, "Track association performance of the best hypotheses search method,"aiuws.unibe.ch. [Online]. Available: http://aiuws.unibe.ch/ccd/publist/data/2013/artproc/TS{_}ESDC2013a.pdf

[11] H. Kurosaki, "Observation of Light Curves of Space Objects," 2009.

[12] H. Rodriquez, K. Abercromby, K. Jarvis, and E. Barker, "Using Light Curves to Characterize Size and Shape of Pseudo-Debris," 2006. [Online]. Available: https://ntrs.nasa.gov/search.jsp?R=20060024592

[13] S. Ostro, R. Connelly, and M. Dorogi, "Convex-profile inversion of asteroid lightcurves: Theory and applications," Icarus, 1988. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0019103588901261

[14] M. Kaasalainen and J. Torppa, "Optimization Methods for Asteroid Lightcurve Inversion I. Shape Determination," Icarus, vol. 153, pp. 24–36, 2001. [Online]. Available: http://www.idealibrary.com

[15] D. Hall, B. Calef, K. Knox, and M. Bolden, "Separating attitude and shape effects for non-resolved objects," The 2007 AMOS, 2007. [Online]. Available: http://amostech.com/TechnicalPapers/2007/NROC/Hall.pdf

[16] M. Sugiyama, Introduction to Statistical Machine Learning. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2016.

[17] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016, http://www.deeplearningbook.org.

[18] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," International Journal of Computer Vision (IJCV), vol. 115, no. 3, pp. 211–252, 2015.
[19] A. Krizhevsky, I. Sutskever, G. H. A. in Neural, and undefined 2012, "Imagenet classification with deep convolutional neural networks," papers.nips.cc. [Online]. Available: http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks
[20] R. Linares and R. Furfaro, "Space object classification using deep convolutional neural networks," in Information Fusion (FUSION), 2016 19th International Conference on. IEEE, 2016, pp. 1140–1146.
[21] B. Jia, K. D. Pham, E. Blasch, Z. Wang, D. Shen, and G. Chen, "Space object classification using deep neural networks," in 2018 IEEE Aerospace Conference. IEEE, 2018, pp. 1–8.
[22] AGI Inc , " Systems Tool Kit (STK) https://www.agi.com/home."
[23] "Space-track.org," https://www.space-track.org/, accessed: 2017-07-31.
[24] R. E. P. W Viehmann, "Ultraviolet and visible brdf data on spacecraft thermal control and optical baffle materials," pp. 0675 – 0675 – 6,1987. [Online]. Available: https://doi.org/10.1117/12.939484
[25] F. Chollet et al., "Keras," https://keras.io, 2015.
[26] "HPC Centers - Non-Allocated Systems." [Online]. Available: https://centers.hpc.mil/systems/nonallocated.html{\#}Hokulea
[27] I. Hussein, C. Roscoe, and M. Wilkins, "Information theoretic criteria for observation-to-observation association," 2014. [Online]. Available: http://www.dtic.mil/docs/citations/ADA616791