Enhancing Cognitive Fusion For Space Situational Awareness

Steven Ingram*, Rupinder Paul Khandpur†, Andrew Zizzi*, Brian Mayer†, Naren Ramakrishnan† and Moses Chan*
*Lockheed Martin Space
Sunnyvale, CA 94089, USA
Email: {steven.ingram,moses.w.chan,andrew.j.zizzi}@lmco.com
†Discovery Analytics Center
Virginia Tech
Arlington, VA 22203, USA
Email: {rupen,bmayer,naren}@cs.vt.edu

Abstract

As the number of objects sent into space and into orbit around the Earth is likely to continue increasing, Space Situational Awareness will continue to require up-to-date methods and algorithms that can predict, and eventually mitigate issues caused by satellite anomalies. Cognitive fusion modeling approach which entails combining informative features from different data modalities presents a unique approach to orbital anomaly detection tasks.

In this paper, we investigate two methods that can potentially make our fusion approach more feature rich and provide higher confidence when estimating if a satellite is "of interest". First, we investigate the use of unsupervised outlier detection algorithms using a set of delta features which represents the change in the value of satellites orbital elements. For this task, we make use of "structured" sensor data that comprises of Two-Line element (TLE) sets which represent satellites orbital path. In the method for unstructured data, we make use of images of satellites in orbit to detect patterns of anomalous behavior. With help of a very comprehensive set of known satellite anomalies as noted by online sources, we will examine and evaluate several scenarios based on historical events that document the utility of these additional techniques for assisting with cognitive fusion modeling.

I. Introduction

With all of the countries, organizations, and space agencies launching satellites into Earths orbit, the task of maintaining Space Situational Awareness is becoming increasingly complex and multi-faceted. Improving methods applied to this task will require up-to-date approaches and algorithms that will be able to predict and prevent satellite anomalies. Many methods already incorporate analysis of Two-Line Elements (TLEs) made available by space surveillance networks, as well as other sources including space-track.org, N2YO [1], and Space News [2]. In some cases, data sources regarding easily predictable craft are more prevalent, such as with the ISS; however, other crafts' data involving behaviors and failures, such as with the Tiangong 1, were insufficient to prevent its uncontrolled descent back to Earth. Probabilistic modeling and data fusion from multiple sources, including text, imaging, radar, TLE, telemetry, could reduce the number of anomalies and time analysts spend reviewing issues during the lifetime of a satellite.

Previous work [1] that used probabilistic modeling and fusion of structured and unstructured data demonstrated an example use case with which a satellite could be classified as "of interest" or not, based on natural language processing markup on recent open-source news information and unprocessed TLE data from Space-Track.org. In a deployed system, "of interest" would imply that an event involving the satellite had occurred, including a maneuver. While this framework showed an improvement in accuracy, precision, and recall metrics when combining open-source text information with TLE data, the model did have some problems generalizing on satellites where no recent open-source source information was available or when a given TLE did not resemble TLEs from the training set. There were a number of satellites that the model was unsure how to label as "of interest" or not due to a lack of features. Increasing the number of features available to factor graph is one method that could reduce the model's uncertainty in whether or not a satellite was "of interest" because it had maneuvered.

This work discusses the use additional approaches to assist with the task of cognitive fusion for satellite anomaly detection in the context of space situational awareness. First, we explore the use of unsupervised methods for orbital anomaly detection using features extracted from TLEs of each satellite including delta (change) of period, semi-major axis and a novel use of satellite's analytical propagation computed using sgp4 [2] procedures. While the unsupervised anomaly detection methods focus on structured data, we also investigate the use of (unstructured) images of satellites produced by sensors at Lockheed Martin's Space Object Tracking (SpOT) facility.

This paper is presented as follows. First, cognitive fusion is briefly described and additional methods are presented as features that could be used to improve that framework. The unsupervised approaches are described and run against real data from events in the past. In addition, images collected from the SpOT facility processed by IBM's image classification tool are presented.

[1] https://www.n2yo.com/
[2] https://spacenews.com/

Finally, results are provided with the aforementioned methods and evaluation metrics are presented with details on how these features could be included in future cognitive fusion frameworks.

II. Problem

To briefly state the cognitive fusion problem, we would like to estimate the most probable world given a set of defined features. The full derivation can be found here [1]. Allow every field in a given database to be a random variable whose domain is the same as the field attribute's domain and let an observed variable X be a deterministic field, and a hidden variable Y be an uncertain field. The hypothesis space, then, of the random variables (X and Y) contain the set of possible worlds because each field has domain equivalent to its respective attribute. Deterministic factors are able to model constraints over sets of variables by outputting 1 if the constraint is satisfied, and 0 if it is violated.For us, we define W to be all the possible worlds with respect to the factor graph's probability distribution η:

$$W = \{\omega \in W_{DB} | \eta_G(\omega)\} \tag{1}$$

Briefly, we can measure observed set of variables, X, to determine the state of the hidden set of variables Y. Having implemented the initial model as solution to estimate the formalized above with Natural Language Processing markup from IBM's Watson Discovery service and unprocessed TLEs, results showed that while the factor graph model showed improvements in classification when combining TLE information with NLP mark-up as opposed to in isolation, it was unable to classify every satellite as "of interest" or not with certainty either below 0.1 or above 0.9. Investigation of the data revealed that in some test cases NLP mark-up was unavailable or TLEs were unlike TLEs from labeled data. Figure 1. shows how many satellites the factor graph was unsure about. We will now present additional sources from which the random variable X, representing the set of all observed variables, could be augmented, thereby improving the graphical model.

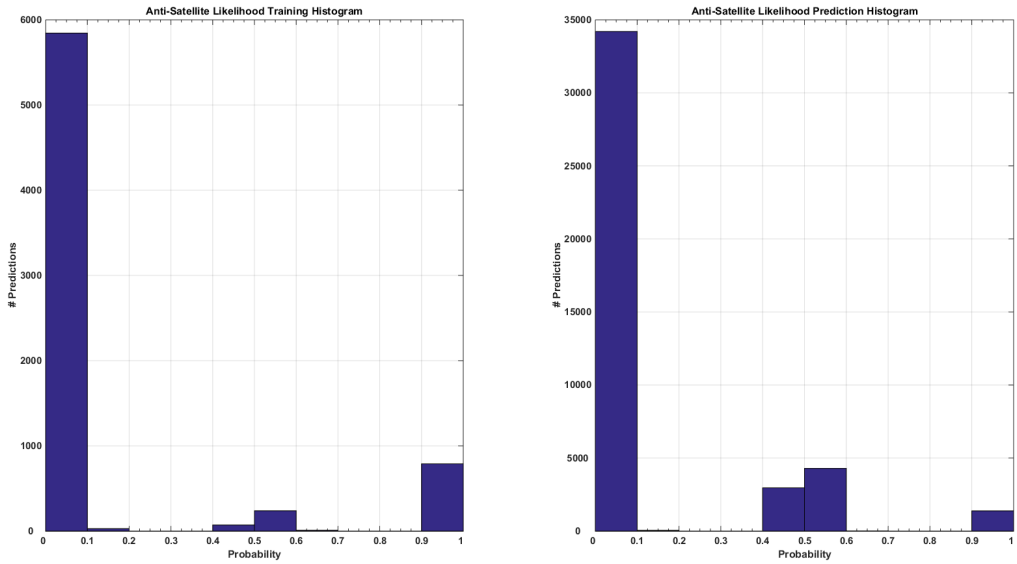

Fig. 1. Results Calibration Plots For "Satellites of Interest".

One method for providing additional context to the results of a factor graph includes building a richer set of features that could supplement the factor graph model to determine the likelihood of something being true or not. In this paper, we present the derivation of additional potential feature candidates, namely the output of unsupervised learning methods applied to TLEs for satellite maneuver estimation and classification of satellites from their images produced from Lockheed Martin's SpOT facility. Where, the context derived from SpOT image analysis provided by IBM's Watson Visual Recognition is the classification of a satellite in the image, which can then be combined with additional information such as what direction the sensor was looking and is expected to be in the image.

The problem is essentially to assign probabilities to the likelihood of specific satellites in an image, assuming that only one satellite is in the image. Multi-class object detection from an image can be solved with deep neural network, and more appropriately, with a convolutional neural network (CNN).

The most fundamental building block of a CNN is the convolutional layer. The following figure illustrates such an example. Equation 2 shows the output of a neuron in a convolutional layer. More in-depth information on convolutional layers can be found at 2.

$$z_{i,j,k} = b_k + \sum_{u=1}^{f_h} \sum_{v=1}^{f_w} \sum_{k'=1}^{f'_n} x_{i',j',k'} \cdot w_{u,v,k',k} \quad with \quad \begin{cases} i' = u \cdot s_h + f_h - 1 \\ j' = v \cdot s_w + f_w - 1 \end{cases} \tag{2}$$

- $z_{i,j,k}$ is the output layer of the neuron located in row i, column j in feature map k of the convolutional layer (layer l).
- s_h and s_w are the vertical and horizontal strides, f_h and f_w are the height and width of the receptive field, and $f_{n'}$ is the number of feature maps in the previous layer (layer $l-1$).
- $x_{i',j',k'}$ is the output of the neuron located in layer $l-1$, row i', column j', feature map k'
- b_k is the bias term for feature map k (in layer l).
- $w_{u,v,k',k}$ is the connection weight between any neuron in feature map k of the layer l and its input located at row u, column v (relative to the neuron's receptive field), and feature map k'

Typically, a rectified linear unit (ReLU) is used after convolutional layers and after fully connected layers. The equation for the operation is shown below with $\mathbf{X}$ being the input and $\mathbf{w}$ and b being the weights and bias, respectively.

$$h_{\mathbf{w},b}(\mathbf{X}) = max(\mathbf{X} \cdot \mathbf{w} + b, 0) \tag{3}$$

Letting $\mathbf{y}$ represent the output of the last ReLU we can now describe the softmax function. The vector $\mathbf{y}$ are the logits and we can turn these numbers into probablities using the softmax function.

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} \tag{4}$$

Next we compute the One-Hot equivalent labels for each image and can compute the cross-entropy for the given input (this is a loss function). Roughly, the cross-entropy measures how inefficient our predictions describe the true measurement.

$$D(S, L) = -\sum_i L_i log(S_i) \tag{5}$$

Now that we have a loss function, we can find all the weights and biases that would give a small value for a correct prediction, and a large value for an incorrect prediction. We want to minimize the following loss function:

$$L = \frac{1}{n} \sum_i D(S(y_i), L_i) \tag{6}$$

This can be done with various gradient descent algorithms which we will not detail here.

With the objective to be minimized defined with labeled training examples, we can now introduce the CNN architecture used to distinguish among two satellites tracked by SpOT. The simplest CNN architecture generally stacks a few convolutional layers (each one generally followed by a ReLU layer), then a pooling layer, and another few convolutional layers (plus ReLU), then another pooling layer, and so on. Typically, the image becomes smaller as it progresses through the network, but it also gets deeper (i.e., with more feature maps) due to the convolutional layers. At the end, a regular feedforward network neural network is added, composed of a number of fully connected layers (+ReLUs), and a final softmax layer to output the prediction. The following section describes the SpOT facility and what CNN architecture was used for satellite classification.

III. Approach

In this section we describe in detail the two approaches we investigate for satellite anomaly detection.

A. Unsupervised Anomaly Detection

In this unsupervised learning approach to orbital anomaly detection, we train two outlier detection classifiers using the two-line mean element (TLE) sets (1) we learn a *robust covariance* estimate [3] to , and thus fits an ellipse to the central data points, ignoring points outside the central mode; (2) using *isolation forest* [4], [5] that comprises of random forest of decision trees, where the objective is to "isolates" observations by randomly selecting a feature and then randomly selecting a split value between the maximum and minimum values of the selected features.

The selected features are derived from each satellite's TLE set which include:

1. **Propagation (change)** is calculated using the Simplified General Perturbations (SGP4 [6]) propagator that is used with TLE sets, considers secular and periodic variations due to Earth oblateness, solar and lunar gravitational effects, gravitational resonance effects and orbital decay using a simple drag model. Such analytical propagators allow fast

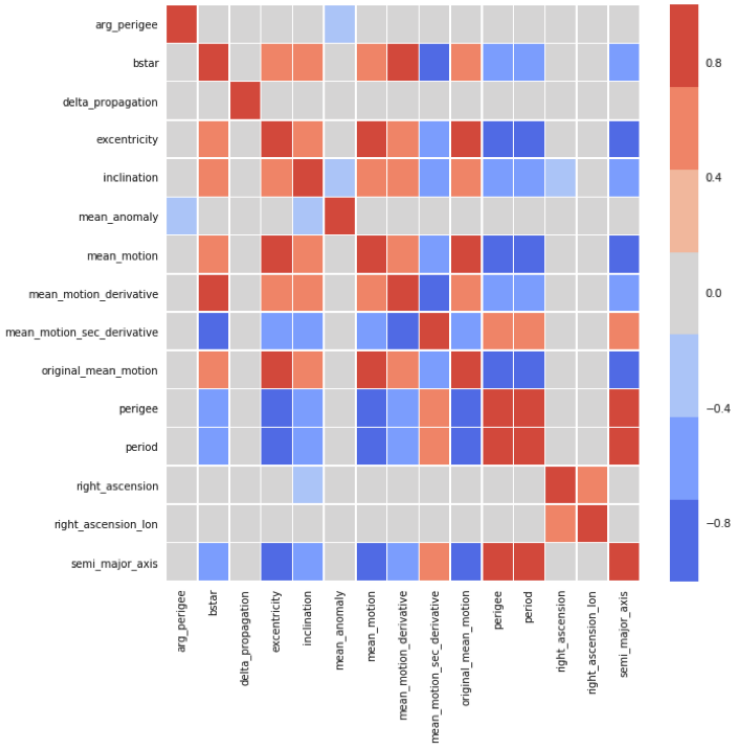

Fig. 2. Correlation matrix of TLE attributes (values from EchoStar 3).

computation of the time-dependent motion of a satellite to provide directly the position and velocity of a satellite at a particular time. The result of which we then use to compute the propagation delta (change) as the Euclidean distance between the actual position (3D point in space) using the timestamp for a given TLE and projected (propagated) position estimated from the preceding recorded, TLE. This feature is not part of the elemental attributes of a TLE set and is unique to this study.

2. **Period (change)** defined as the time taken by the satellite to complete one orbit around the earth.
3. **Semi-major Axis (change)** defined as the average distance of the satellite from the Earths center, thus describes the size of the satellites orbit. Similar to computation of Delta Propagation when using Period and Semi-major Axis we use their (change) delta values in consecutive TLEs. Each of the features are then scaled by removing the mean and scaling to unit variance before fitting the outlier detection classifier.

The correlation matrix (shown in Figure. 2) of TLE features provides a guide for selecting specific feature sets for training the classifiers.

B. SpOT

To facilitate identification and tracking of satellites for the purposes of Space Situational Awareness, Lockheed Martin has stood up multiple sensor facilities in various locations, both in the US and Australia. As part of this sensor network, the SpOT facility is a reconfigurable telescope array located at Lockheed Martins Santa Cruz Test Facility (SCTF) in California. Three 1-meter telescopes are mounted on motorized carts on a three rail lines: North-South, Northwest-Southwest, and Northeast-Southwest. The telescopes are housed inside Astrohaven 12 foot domes for environmental control and an unobscured sky view. Figure 3 shows an elevated view of the facility.

Each telescope is an afocal, three-mirror, on-axis anastigmat, with glass optical surfaces. A modest quality Basler CCD camera is used to acquire data in the 1-meter optical path, as well as from the 106mm side scope (Takahashi FSQ-106). R&D at SpOT investigates and demonstrates multi-phenomenology data products: speckle cameras for closely spaced objects, high speed cameras for hypertemporal imaging, polarimetry for object identification and characterization, spectrometry for target identification, characterization, and filter wheels for color photometry. It has sensing capability for visible through mid-IR range (450nm - 8μm).

Given the images like those generated by SpOT, and without any additional context, a human would have difficulty discerning between one satellite and another. To assist with satellite classification, and by extension, provide additional context to a given satellite for space situational awareness, a CNN was implemented and trained with IBM's Visual Recognition service. The

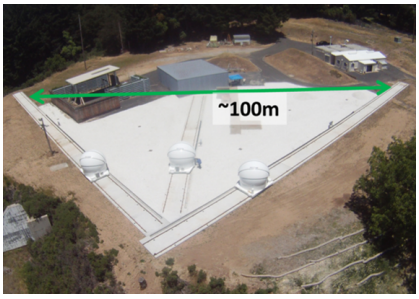

Fig. 3. Elevated View of SpOT Facility Showing Three-Rail System

Fig. 4. Elevated View of SpOT Facility Showing Three-Rail System

architecture of the CNN is broken up into several layer groups shown in the figures III-B, 6, 7, III-B, and 9. The main sections of the architecture consist of Input, Hidden, Softmax, and Output layer groups. Input images are initially down-scaled in the ImageScaler and then passed to 26 hidden layer groups consisting of Convolution, Batch Normalization, Multiplication, Addition, and Activation operations. Figure III-B shows the image scaler at the start of the neural network. Figure 6 illustrates one example of these 26 hidden layer groups that immediately follow the ImageScaler operation.

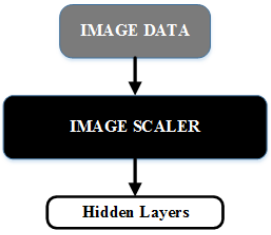

Fig. 5. Input Layers

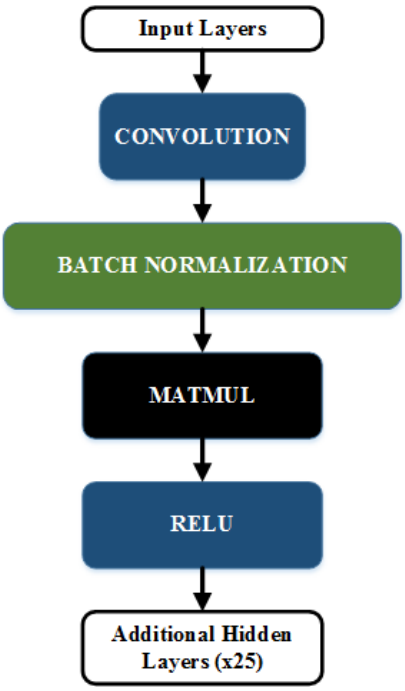

Fig. 6. Hidden Layers (1 of 26)

After the last final Relu operation, at the end of the 26th hidden layer group, an average pooling operation followed by a final convolutional operation connects to the softmax layer where the result is finally sent to the output layers. Figure 7 shows the pooling layer, convolution layer, and softmax layer.

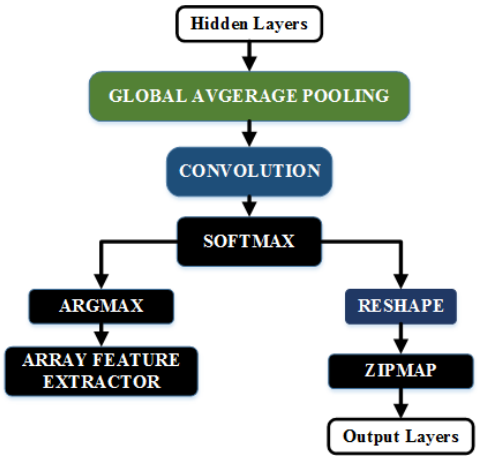

Fig. 7. Pooling, Convolution, and Softmax Layers of CNN for Satellite Classification

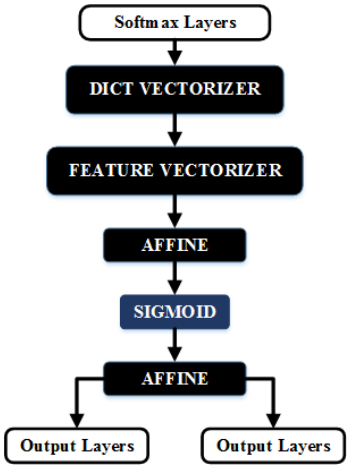

Fig. 8. Output Layers

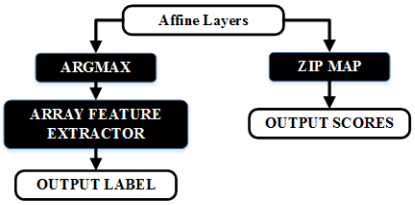

Fig. 9. Label and Scoring Layers

The final output operations consist of an affine transformation, a sigmoid activation, and a second affine operation. The affine operations takes on input data (Tensor) and produces one output data (Tensor) where the affine function, $y = alpha * x + beta$, is applied to the tensor elementwise. The sigmoid function is defined as y = 1/(1+exp(-x)), which is applied to elementwise. Finally, the scores for each satellite class are output for the given image. The scores sum to one, with the largest value associated with the most likely object in the input.

The initial model provided by IBM's Visual Recognition contains weights and biases for the input layers and subsequent hidden layer groups that were learned for other image classification tasks, and the layers nearest the output were left to be retrained using our satellite images. This allowed training to complete without the number of training examples that would have been required if the entire network would need to be trained from scratch. This reduced our training requirements by at least one order of magnitude. The model was trained with 800 images containing "SAT_01" and with 1100 images containing "SAT_03." These images are of the file type .png with dimensionx 776 pixels wide by 582 pixels long. Results from the test set are explained in the next section.

IV. Experiments

Ground-truth Dataset of Satellite Anomalies: We were able to identify three open sources of information on satellite anomalies (listed in Table. I) from which a total of 418 (anomaly) events were collected that included 298 unique satellites. One of the primary challenges that we face is the lack of availability of such datasets which are open to researchers (also highlighted in a 2014 study [7] done by RAND Corporation). Still, to the best of our knowledge, this is most extensive satellite anomaly dataset compiled for publication.

Satellite Catalog & TLE-ETL Pipeline: Datasets pertaining to satellites are divided into two parts (1) A complete catalog of 43,064 satellites containing their high-level descriptors (2) TLE orbital paths for 4242 satellites that amount to over 18 million unique TLE sets sourced from an online source (Space-Track.org [3]). The TLE dataset also includes all satellites that are reported in anomaly dataset. Also, a comprehensive data pipeline (available in source code) has been implemented to load, transform and query TLE information which are then used for data exploration and model building tasks.

TABLE I
SUMMARY OF SATELLITE ANOMALY DATASET.

Source	Num. of Anomalies	Description
Satellite News Digest [4]	266	On-Orbit Failures / Reported Outages
NOAA [5]	98	On-Orbit Failures / Reported Outages
JPL [6]	47	Power Subsystem Failures
Manual Searches	7	On-Orbit Failures / Reported Outages

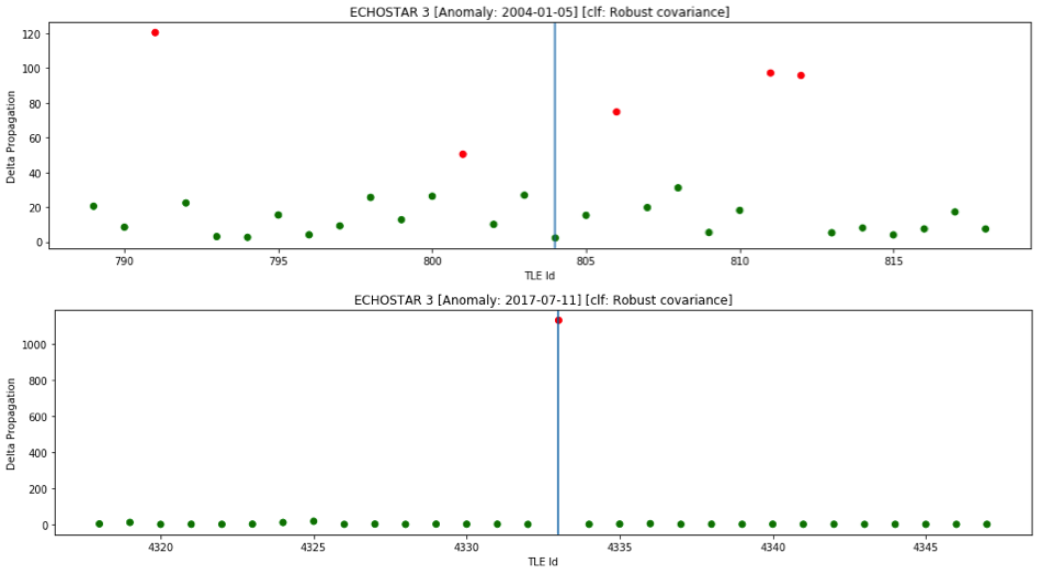

Fig. 10. A scatter plot showing detected outliers (red points) using delta propagation values of satellite EchoStar 3, for which the reported anomaly events are annotated as vertical spans (in blue).

Orbital Anomaly Detection Performance: The performance metric chosen for the experiments is true positive rate (TPR), since we cannot validate all the detected anomalies with a verified ground truth source, we calculate the TPR conditioned only known true labels. This gives us a good picture of the performance of the detection methods and the usefulness of the selected features. Further, we noticed from scatter plots of detected outliers and inliers for different satellites, that in many cases anomalous behavior was detected either before or after the reported date of anomaly (see Figure. 10). This reporting lag could be attributed to when it is officially reported via official sources. For this reason, we plot the TPR over varying window sizes. Each window size is defined as number of adjacent TLEs selected to right & left (+/-) of the specific TLE for which there is a known anomaly. So, if there is a detected anomaly in that specified window we count that as true positive. We report the results of outlier detection in Figure. 11. The plots show similar performance for both classifiers. We notice a sharp jump in TPR from .33 to .6 when we increase the window size to (+/-) 1 TLEs for classifiers trained on delta propagation values which also outperforms other feature selections used.

SpOT Analysis: The table below shows the results from the test set of images the neural network processed for satellite classification. Once the model had been trained, validation of the model was performed with a test set of images that were held out during the training of the network. The test set contained images 2177 images of "Sat_01", 1135 images from one day and 1085 images from another day. There were a total of 2230 test images of "Sat_03", 1178 images from one day and 1052 images from another day. Actual satellite names and associated dates are not reported for public information release purposes. The table below summarizes the inference results from the test set. From this, we have a test comprised of 4677 images. 2200 images actually of Satellite 03 were correctly classified as Satellite 03 and 2016 images of Satellite 01 were correctly classified

[3]https://www.space-track.org/documentation

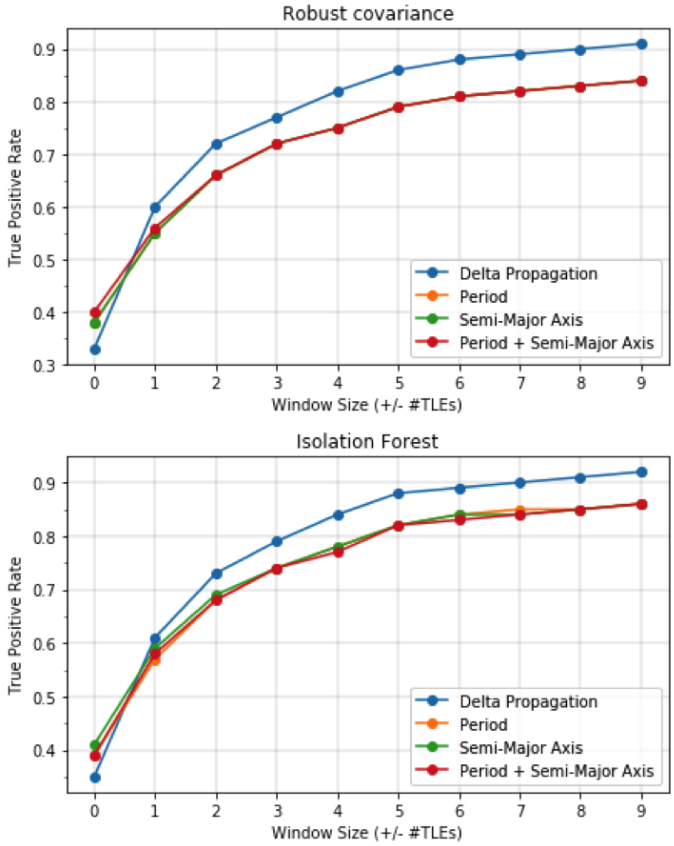

Fig. 11. Probability of anomaly detection over varying window sizes for different features.

as Satellite 01. There were 161 images where Satellite 01 was incorrectly classified as Satellite 03 and there were 300 images of Satellite 03 that were incorrectly classified as Satellite 01.

n=4677	**Predicted: SAT_01**	**Predicted: SAT_03**	
Actual: SAT_01	2016	161	2177
Actual: SAT_03	300	2200	2500
	2316	2361	

Fig. 12. Confusion Matrix for SpOT Satellite Classification

From this, we derive the model's overall accuracy, precision, recall, and F1 score. Accuracy is calculated as $(True\ Positive + True\ Negative)/total = (2016 + 2200)/4677 = 0.957$. For Precision, we calculate as $True\ Positive/(True\ Positive + False\ Positive) = 2200/(161+2200) = 0.959$. Recall is calculated as $True\ Positive/(True\ Positive + False\ Negative) = 2200/(2200 + 300) = 0.956$. Using recall and precision calculated earlier, we can compute F1 Score as $2 \cdot \frac{precision \cdot recall}{precision + recall} = 2 \cdot \frac{0.959 \cdot 0.956}{0.959 + 0.956} = 0.957$. The results of these metrics are summarized below:

METRIC	**VALUE**
ACCURACY	0.957
PRECISION	0.959
RECALL / SPECIFITY	0.956
F1 SCORE	0.957

Fig. 13. Metrics for SpOT Satellite Classification

V. Conclusion

In this paper, we have investigated new approaches to satellite anomaly detection using unsupervised learning techniques that rely on features extracted from (structured) two-line-element (TLE) sets of satellite and analysis of images of satellites generated from Lockhead Martin's SpOT facility. We have discussed the results of out exploratory analysis and have shown that they can be quite effective for the anomaly detection task. Together, these approaches allow a more rich feature set. Some of the key learning points from our analysis include:

1) Not all anomaly types are appropriate for training satellite orbital anomaly detection models (using TLEs) since they might pertain to failures or outages of satellites on-board systems which may or may not impact its orbit.
2) Robust Covariance method is suited to low-dimensional data whereas Isolation Forest is an efficient way to detect outlier even in high-dimensional settings. Further, new experiments can be performed that explore are more expansive (TLE attributes) feature sets.
3) SGP4 method used to calculate positions (delta propagation) based on TLEs is applicable to only Low-Earth orbit (LEO) satellites. Hence, any non-LEO satellite anomalies should not be used with the classifier that uses delta propagation.
4) A pre-trained convolutional neural network architecture provided by IBM's Visual Recognition Service can be modified to perform satellite classification based on images captured by SpOT.
5) The CNN was able to achieve notable results, achieving an F1 score of 0.957.

Future work would include training multiple factor graphs with various weights and mixtures of these proposed features to determine which model is most optimal making determinations about satellites for space situational awareness.

References

[1] S. Ingram, M. Shaw, and M. Chan, "Applying cognitive fusion to space situational awareness," in *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, Maui, Hawaii, 2017.

[2] D. A. Vallado, P. Crawford, R. Hujsak, and T. Kelso, "Implementing the revised sgp4 in stk," in *Proceedings of the AGI User Exchange*, 2006.

[3] P. J. Rousseeuw and K. V. Driessen, "A fast algorithm for the minimum covariance determinant estimator," *Technometrics*, vol. 41, no. 3, pp. 212–223, 1999.

[4] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 2008, pp. 413–422.

[5] F. T. Liu and Z.-H. Ting, Kai Ming & Zhou, "Isolation-based anomaly detection," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 6, no. 1, p. 3, 2012.

[6] F. R. Hoots and R. L. Roehrich, "Models for propagation of norad element sets," Aerospace Defense Command Peterson, Office of Astrodynamics, Tech. Rep., 1980.

[7] D. A. Galvan, B. Hemenway, I. Welser, D. Baiocchi *et al.*, "Satellite anomalies: Benefits of a centralized anomaly database and methods for securely sharing information among satellite operators," RAND National Defense Research Institute, Santa Monica, California, Tech. Rep., 2014.