Machine Learning for E-O Data and Imagery Event Detection

John Ebeling, **Duane DeSieno**, **Jacob Hansen**, **Christopher Tschan**
Data Fusion & Neural Networks
Carolyn Sheaff
Air Force Research Laboratory/RIED

ABSTRACT

As space grows ever more congested, achieving and maintaining Space Domain Awareness has been the focus of many research efforts. Electro-Optical (E-O) space surveillance observations are growing in importance and quantity. This paper presents a prototype implementation of a modular pipeline that characterizes space situations by leveraging both electro-optical data and imagery. The pipeline is designed to be deployed at the space telescope sensor location in order to immediately process and autonomously alert operators to data and images that contain interesting events, such as multiple satellites in proximity. Ultimately, this algorithm would also provide automated near-real time imagery annotation significantly reducing time to respond to emerging space events. This allows individual E-O sensors to immediately contribute to Space Domain Awareness and opens new opportunities for higher level fusion to be performed on centralized data processing platforms.

This work culminates two Air Force Research Laboratory (AFRL) projects, the Rapid Discovery of Evasive Satellite Behavior (RDESB) and Electro-optical Pre-custody Threat Warning (EPTW). The current iteration of the prototype processes a single E-O image in under 2.5 seconds. The pipeline prototype was first trained and tested using simulated E-O data from the SJ-21 (satellite catalog object 49330) and Compass G2 (satellite catalog object 34779) close approach and grappling scenario in January of 2022. As the objects' visual separation decreased the pipeline correctly identified and annotated a visual proximity event. A Focus of Attention algorithm was used to identify real data of interest from the Unified Data Library for additional testing. Using this data, the pipeline prototype produced similar promising results.

The prototype pipeline utilizes multiple sets of Neural Networks (NNs) with Machine Learning and Data Fusion algorithms to extract, detect, and categorize scenarios. Convolutional, traditional feed forward, and autoencoder NNs are all employed for object extraction, light curve analysis, and anomaly detection respectively. We investigate Gaussian Mixture Model fitting as a fast solution to source extraction. The Hungarian algorithm is used to solve the assignment problem. A heuristic-based clustering approach is used to classify time series of error scores. This paper also discusses plans to replace the heuristic-based clustering with NN-based clustering as more annotated event timelines become available. Imagery and derived data analysis subprocesses run in parallel pipelines with their outputs being fused to produce an event timeline that can be resampled to an operator's desired scale. Recommendations for future optimization and research are also presented.

1. Background

As space grows ever more congested and contested more Electro-Optical (E-O) data continues to be produced. This virtual flood of E-O data, mostly from ground-based equipment, seems to be on track for continual growth as more satellites launch and as the technology to observe becomes cheaper. The unhindered increase in data supply demands an increase in processing algorithm efficiency to contextualize E-O data and quantify its importance with specific focus as to whether or not to bring specific observations to a Space Domain Awareness operator's attention.

Much work has been done to related to applying Neural Network (NN) technologies to object localization and extraction. Fletcher et al (2019) used Convolutional Neural Networks (CNNs) for object extraction [1]. This approach has been extended to using ground-based multi-spectral imagery and star localization for low field-of-view

scopes to great effect (see Gazak et al 2023) [2]. Lucas et al (2019) [3] produced automated resolution scoring using CNNs for LE-O image scoring. Woodward et al (2021) also developed CNNs for RSO identification in the GE-O regime [4]. We employ similar techniques, combining CNNs further with Autoencoder style NNs to learn normal image patterns (on a satellite-sensor pair basis) as a primary means of anomaly detection.

A second approach to E-O data-based anomaly detection uses light curve analysis rather than image analysis to detect changes. Meredith et all (2023) showed that as Intelsat 10-02 and MEV2 approached closely in visual proximity, the individual light curves eventually melded into a new light curve with significant differences from the original (although dominated by the shape of the brighter object). This enabled them to estimate the celestial positions of the combined objects during glint events using the Source Extractor algorithm and centroid estimation using non-resolved imagery [5]. Our approach similarly relies on the idea of the melding of two light curves creating a new blended curve. A feed-forward NN is trained on the nominal light curve for the object of focus, and outputs from this network may be processed in parallel to produce secondary evidence for event detection.

Shakun et al (2023) compared two pipelines built for similar purposes: the Ukrainian LEMUR and Polish Poznań Satellite Software Tools (PSST) image processing pipelines [6]. Similarly, Stewart et al (2023) presented work on a pipeline for space-based sensors to produce Space Situational Awareness (SSA) information [7]. Our approach differs from these works mainly in scope. The outcome of the cited works was to calculate to sub-pixel accuracy the position of a Resident Space Object (RSO) based on E-O observation (both imagery and extracted visual magnitude / celestial position). Our pipeline steps beyond RSO localization with a focus on quantifying the abnormality of an observation given historical context. In a similar vain, Patil et al (2021) designed a pipeline for asteroid collision detection using E-O observational data with CNNs, but our work generalizes beyond proximity detection to anomalous proximity detection [8].

In summary, much research has been done to apply CNNs to object localization within imagery for many satellites at various regimes. In parallel, research also shows that shifting light curves can be used for melded object discrimination and detection. This work focuses on applying these well-known techniques to the abnormal event detection and classification problem using Autoencoder NNs as the primary means of quantifying abnormality. Our approach utilizes two parallel pipelines for processing imagery separately from analytical E-O data (Right Ascension, Declination, Solar Phase Angle, Visual Magnitude). The results from each are fused to create an event timeline for the possibly unresolved or pre-custody object of study. The entire pipeline can operate on outputs from either or both parallel pipelines. The result of our work is a prototype system that uses E-O data and imagery for Pre-custody Threat Warning (EPTW) generation and the Rapid Discovery of Evasive Satellite Behavior (RDESB).

2. Methodology

The problem that the EPTW System seeks to solve is ‘to provide “an automated process that runs on-site at sensor locations to immediately recognize potential threat events in imagery to alert operators on relevant timelines” (e.g., <10 seconds) to include abnormal event detection scores and event characterization with confidences plus trust scores. This allows a human operator or an automated system the option of collecting more observations of “interesting” events as they develop.

A simplified, top-level data flow plan for the algorithms actually processing E-O imagery and E-O data in the EPTW system is displayed in Fig. 1. A sensor produces one of or both E-O data and E-O imagery at the top of the figure and the output of the EPTW system is represented in the display box at the bottom. There are two vertical paths shown, one for E-O data (left) and the other for E-O imagery (right). The heritage of the modules in Fig. 1 is depicted by color. The blue modules, particularly on the right vertical path were created during the SBIR-funded EPTW phase I projected at the prototype level and include some that are planned for development later stages. The purple module on the left vertical path were derived from DFNN’s RDESB phase II contract with AFRL.

An overview of the image processing data flow is as follows: the image (usually expected in FITS format) enters the Extract module, where the image is normalized, objects of interest are identified and localized using a CNN

(YOLOv5 in our implementation), and additional analytic statistics are calculated for each object and for the image at large. The Detect FITS module takes the outputs of the Extract module and runs algorithms for general anomaly detection, performs data association with tracks stored in a local catalog by solving the assignment problem using the well-known Hungarian algorithm, and passes these as inputs into a pre-trained Autoencoder to detect anomalous activity. The output of the Detect FITS module consists of the input observation augmented with statistics from the Extract pipeline and normalized error scores from the Autoencoder.

The independent E-O Data processing flow is as follows: the datum (usually containing a minimum of solar equatorial phase angle and visual magnitude at a measurement time) enters the Detect E-O module. Here, a feed forward NN that was pre-trained to capture the normal light curve of the object of interest is used to generate a residual visual magnitude for the given datum. The datum, accompanied with this residual, is fed into an Autoencoder which was pre-trained on the normal distribution of light curve datums and their residuals, generating error scores. The output of the Detect E-O module consists of the input observation augmented with normalized error scores from the Autoencoder.

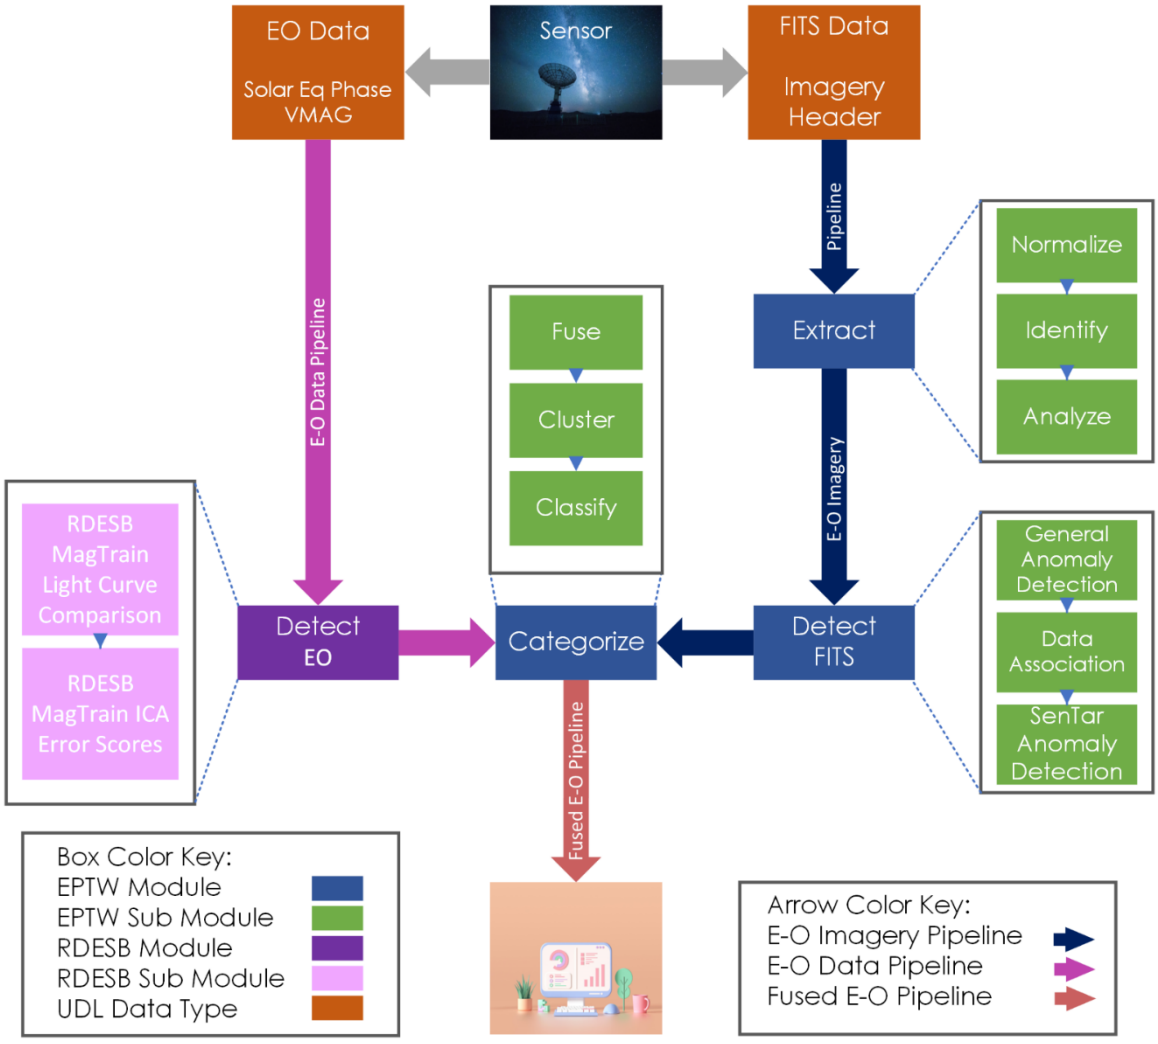

Fig. 1. Top-level Data and Processing Flow diagram for EPTW phase II.

The outputs of the Detect E-O and Detect FITS module are first classified in the error space in sequences of 3 datums with first and second differencing applied to capture temporal error activity to produce a 6-dimensional

vector. The classification in its current implementation is rule-based, using the differences and overall error values to bound areas in 6-dimensional space with hyperplanes to heuristically classify the situation. The hope is that with more label situational E-O data in the future this can be replaced by a feedforward NN. A simple clustering model is trained on several examples of data points and labels to create situational datum clusters within the space. New data, once labeled, then is added to each cluster. The clusters are dynamic and expand their radii to fit the point. The distance of the new point to the center is used as an imprecise metric for similarity to any training data (ie this would be a way to identify if something had been left out of the training set). The data is then resampled at a given frequency so as to bin and match error detections within the same timeframe. It then aligns the most recent resampled data frames from both parallel pipelines in time and combines or fuses their error scores them using specified weights. The result is an event timeline for an object based on that sensor's data which categorizes observations as they come in, quantifying their abnormality, classifying the threat situation, and quantifying the uncertainty of the classification.

2.1 Test Scenario

We will refer to a sequence of proximity events and an eventual grappling event as we evaluate the performance of each module in the first-ever built prototype of the EPTW system. The sequence we used is the interaction between Chinese satellites SJ-21 (49330) and Compass G2 (34779) in January 2022. A basic narrative of the events can be constructed from Fig. 2 showing the propagated normed distance between the two satellites through time from relevant TLEs.

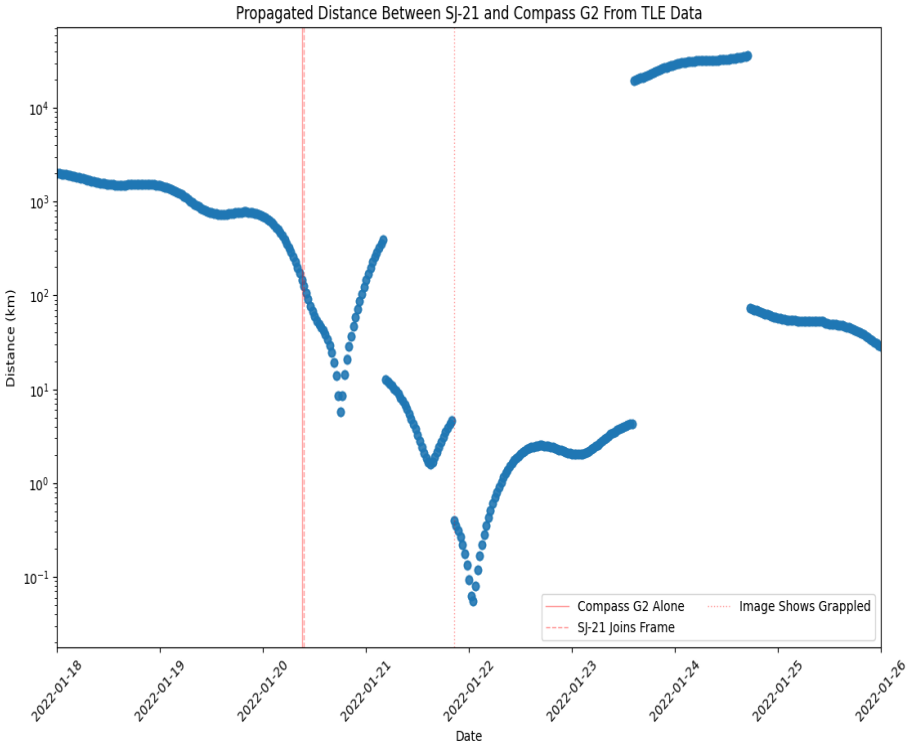

Fig. 2. The Euclidean distance between SJ-21 (49330) and Compass G2 (34779) over time. Distance is shown on a log scale and was recorded in km. Vertical lines mark timepoints when, in the simulated imagery, SJ-21 joins the frame and when it visually appears the satellites have grappled.

We use the data from a single sensor (EXO1488) and our approach focuses on protecting the target, rather than following the chaser. We believe this approach is more operationally oriented because pre-custody threat warning

assumes custody of the focus satellite to be protected, and not necessitating custody of the threatening satellite (in this case SJ-21).

Fig. 2 shows a series of proximity events in the distance, the lowest of which shows the distance reaching well below 1 km. Though we don't observe a flat line of propagated distance at zero in any portion of the graph, we believe this is due to a lack of granularity in the TLE data. The imagery was automatically labelled as an apparent grappling event based when the two objects became indistinguishably close.

Due to a lack of real imagery for this event from government sources on the UDL we simulated imagery using the DF&NN Imagery Generating Simulator (DIGS). One important feature we hope to develop in the future includes the simulation of light curves. As this has yet to be developed, we used real world data from the EXO 1488 sensor for our evaluations.

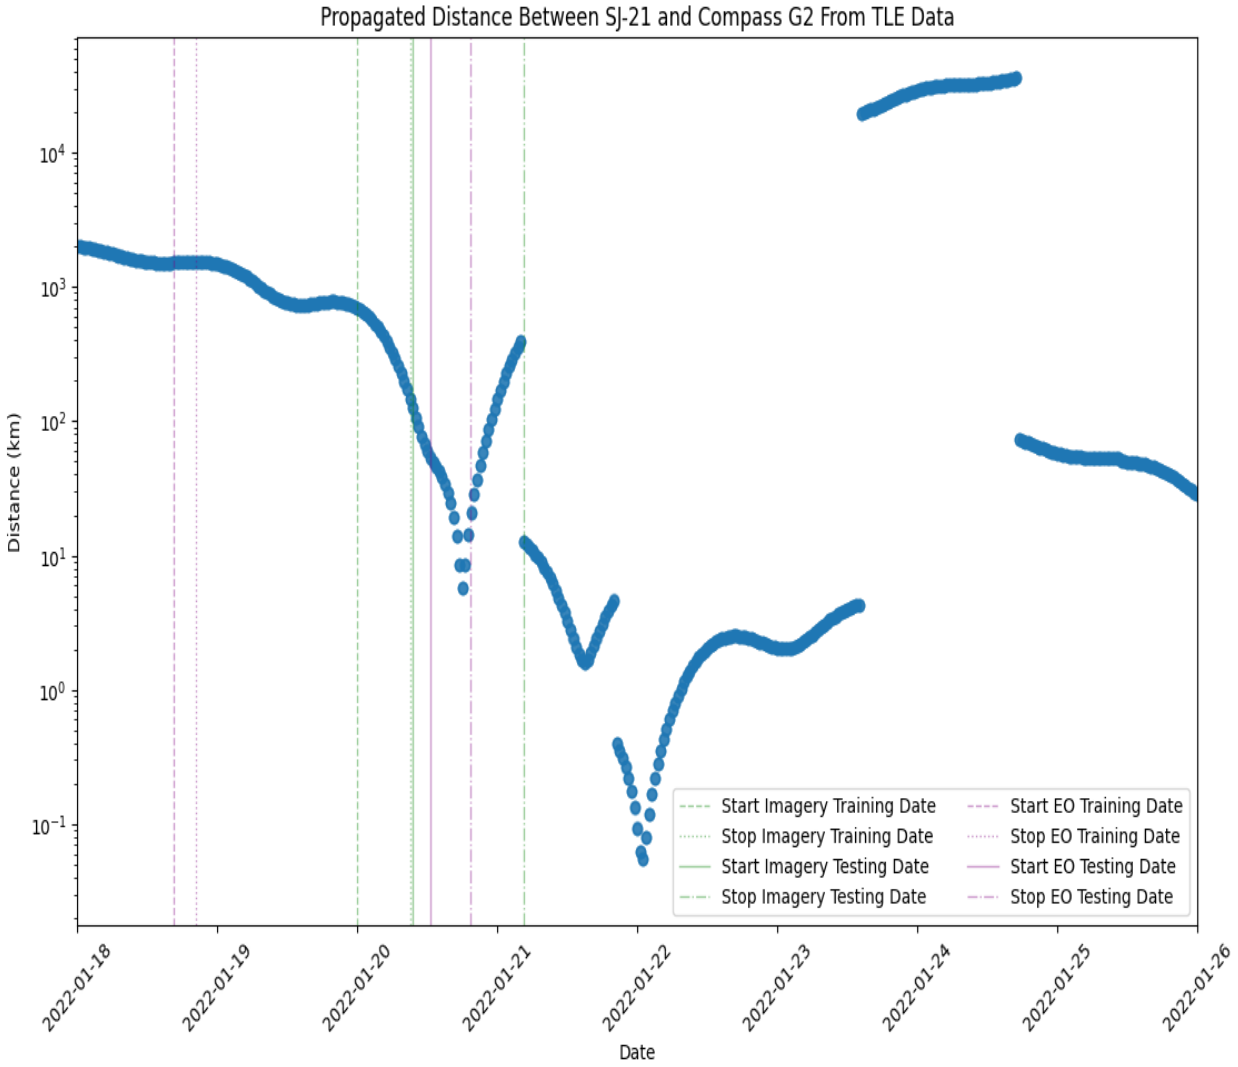

Fig. 3. The Euclidean distance between SJ-21 (49330) and Compass G2 (34779) over time shown in Fig. 2, but now overlayed with vertical lines marking the training and testing (start-stop) times for each data source (E-O and Imagery).

Because there were two E-O separate data sources (time series of E-O data and E-O images), we show separate training and testing periods for each data source (Fig. 3).

3. Results

In this section of the report, we present the module level evaluation of the EPTW system. First, we establish module evaluation criterion, with an explanation for each criteria's selection. Then we summarize each module's performance and break into subsections for each module.

3.1 Module Level Evaluation Criteria

To begin, we review our common criteria for evaluating the module performance, as shown in Table 1. We've selected four criteria as the basis of evaluating each module. The target values listed give DF&NN's goal for performance by the end of future EPTW development.

The first is "Runtime" which was included for the purpose of satisfying a constraint to have a process that runs at sensor locations in under 10 seconds. Since there are five modules running in series (though some should run in parallel in actual deployment of the system) we've assigned a maximum of two seconds for each module to process an image. Specifically, the "Runtime" is the mean time in seconds it takes to process a single image ± the standard deviation over 100 runs. Note that there is a slight chance of increase in runtime given the number of objects in the image, but we use an image with two objects from the SJ-21 Grappling Scenario for evaluation.

The second is "Max RAM Usage". This constraint forces our development to use a reasonable spatial complexity when operating, ensuring we don't run afoul of exorbitant hardware costs. With 16 GB being the RAM space found on many laptops in industry, we hope to keep hardware costs minimal upon deployment.

Table 1: An overview of the criterion used to evaluate each module. The five criteria listed are Runtime, Max RAM Usage, Precision, Recall, and TRL. The "Units" column states the unit of measurement for each criterion. The "Purpose" column gives the stated purpose of the associated criteria. The "Target Values" column define the end of Phase II target value for each criterion.

Criteria	Units	Purpose	Target Values
Runtime	Seconds (s)	To keep the entire processing time within the 10 second time constraint (ie, low temporal complexity)	<= 2 s
Max RAM Usage	Gigabytes (GB)	To keep hardware costs to a minimum (ie, low spatial complexity)	<= 2 GB
Precision	Unitless	To ensure a high precision / low false positive rate (ie, the system didn't give too many false alerts)	>= 0.95
Recall	Unitless	To ensure a high recall / low false negative rate (ie, the system failed to detect a very small number of events)	>= 0.99

The third is "Precision". In machine learning literature, precision is measured as the number of true positives (TP) over the sum of the number of true positives and the number of false positives (FP). Thus, requiring a high precision of 0.95 or greater ensures that operators won't be bothered by too many false positive alerts upon deployment.

The fourth is "Recall". In machine learning literature, recall is measured as the number of true positives over the sum of the number of true positives and the number of false negatives (FN). Thus, requiring a high recall of 0.99 or greater ensures the system and operators don't miss events.

The performance of each module based on these criteria for the SJ-21 and Compass G2 grappling scenario is given below in Table 2. The colors reflect the criteria's performance in relation to their target values. Green means acceptable, yellow indicates a focus area for improvement, and red signifies a value well below the target that should urgently be improved in future EPTW development.

Note that all testing was completed on a laptop with a 12th Gen Intel(R) Core(TM) i7-12650H (16 CPUs), ~2.3GHz, and 16GB of RAM. We should also note that though this is a module level evaluation, running the entire prototype system in series for a single image and E-O datum had a runtime of 2.167 s ± 0.123 s with RAM usage of a maximum of 1.183 GB. The lower maximum RAM usage for the full run is slightly smaller than that of running the Extract module individually and may be related to some optimization inherent in the laptop processor. However, we can with confidence state that our prototype system meets the objective of processing an image in under 10 seconds with ease. We offer comments on each module's performance following Table 2 in the appropriately named subsections.

Table 2: The evaluation of each module of the EPTW prototype system in each of the applicable criterion.

Module Name	**Runtime** $\mu \pm \sigma$	**Max RAM Usage**	**Precision**	**Recall**
Extract	1.543 s ± 0.087 s	1.226 GB	0.993	0.712
Detect FITS	0.021 s ± 0.003 s	0.462 GB	0.971	0.324
Detect E-O	0.036 s ± 0.004 s	0.493 GB	Insufficient Data	Insufficient Data
Classify	0.470 s ± 0.021 s	0.496 GB	0.445	0.348

3.2 Extract Module Performance

The Extract Module's precision was on target at 0.993. In other words, only a small number of detections from the YOLO model didn't correspond to real objects in the frame. The recall of 0.712 is more concerning however, as that shows that a significant portion of objects within the frames weren't found by the YOLO model. One large contribution to the low recall is the observed limitation of our trained model to detect two separate objects within roughly 13 pixels of each other. There are however many simple ways to improve performance in future work including retraining on a less noisy dataset and exploring other models besides YOLO to perform object detection and classification. We've experimented with both in initial EPTW development and are excited to continue efforts to improve the Extract Module's performance in future work.

3.3 Detect FITS Module Performance

The Detect FITS Module's precision was also on target at 0.971. The precision and recall for this module were calculated as error scores above some threshold being anomalous and error scores below that threshold being normal. The high precision shows that only a small number of data points were incorrectly labeled as anomalous. The extremely low recall of 0.324 is most likely due to the patchy performance of the Extract Module. In many cases the Extract Module missed both objects in a sequence of frames and this effect carries through to the Detect FITS module, so that it lacks data for many frames. Comparing a perfect classification in the top of Fig. 4. to the one with recall of 0.324 in the bottom of Fig. 4. shows that despite the low recall, the Detect FITS Module was still able to capture many of the early sequences of error shapes. And as the EPTW System's objective is about producing reliable and predictive threat warning, it is the early signs that are of paramount importance.

3.4 Detect E-O Module Performance

Recalling our previous discussion on using Compass G2 as our focus satellite, this approach relies on Compass G2 receiving a high quantity of observations from a single sensor to produce sufficient data for analysis. Taking the narrative from TLE data as ground truth, Compass G2 had little available data to link signatures in the error space with actual orbital changes. We thus refrained from presenting precision and recall because of the scarcity of the E-

O data. In similar scenarios for real-time deployment, we remind the reader of the Categorize module's ability to operate on a single source of information (E-O data or E-O imagery). The Detect E-O Module's runtime with the data we did have was well within the target parameters as well as the RAM usage.

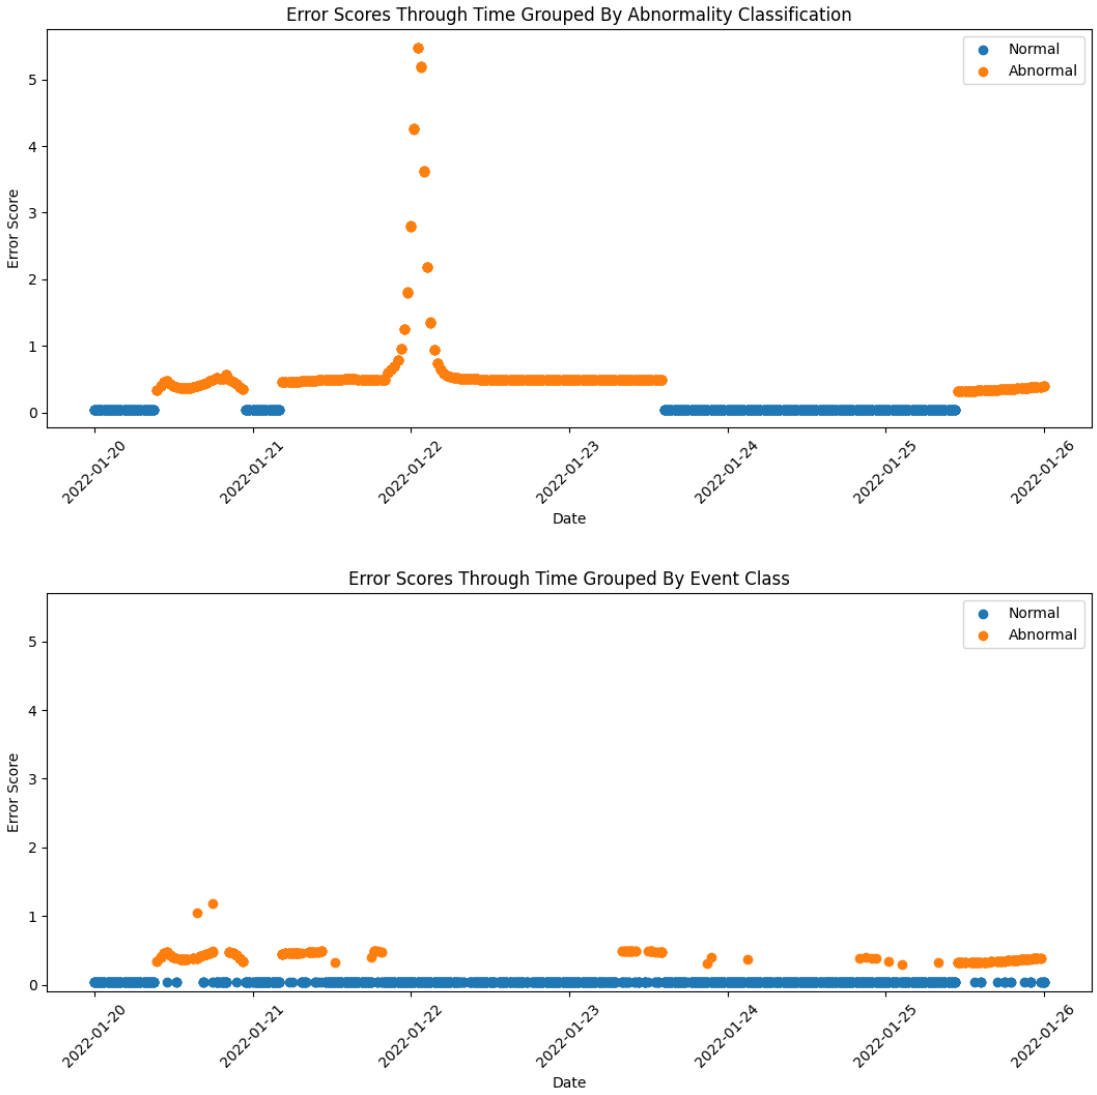

Fig. 4. The Detect FITS Module output ICA Error scores grouped as abnormal and normal. The top shows a perfect run, and the bottom shows the output when the output of the Extract Module is used as input to the Detect FITS Module using the test scenario data.

3.5 Categorize Module Performance

As demonstrated in Table 2, the prior Extract and Detect FITS modules missed a portion of data points, with either the object detection NN either missing the object, or the object's state being mischaracterized by the SenTarICA Errors. The Categorize Module achieved a precision of 0.445 and a recall of 0.348. This however is slightly unfair since many of its misclassifications were caused by classifying something that was an Approach or Backoff as a Proximity event. However, the error scores still indicate a pre-custody threat, and though we intend to improve accuracy in the future, this shows that the prototype capability of the Categorize Module is still of assistance to the

material operator. We seek to show that despite these increased errors the solution still functions as a pre-custody threat warning system, though with elevated uncertainty.

Fig. 5 shows the resulting event timeline, with colors indicating different classifications for each datum. The banding of different colors is the key to visually narrating the situation. Bands of Approach and Backoff events appear in succession in Fig. 5, reducing the likelihood of a collision or fly by event and increasing the likelihood of a rendezvous, grappling, or docking event.

In short, despite the relatively low performance in precision and recall, the Categorize module still provides a meaningful situation estimation to the operator. Notice that these first warnings (oscillating bands of color) are given almost a day before the grappling event, giving the disadvantaged operator plenty of time to react and plan an appropriate Course of Action (COA) including increasing the surveillance on the target to improve the resolution of the EPTW system.

To summarize our results, the Categorize Module in the EPTW proof-of-concept system was demonstrated to provide early warning of the pre-custody threat constituted by SJ-21. That stated, there is much room for improvement. This SJ-21 and Compass G2 sequence was the singular grappling event we generated simulated imagery for to show proof-of-concept functionality for this initial phase of development. We look forward to the opportunity to collaborate with space operators in designing useful displays, creating a larger labeled database, improving precision and recall with Characterization and Trust NNs, and further assessment of the Categorize Module in future version of the EPTW System.

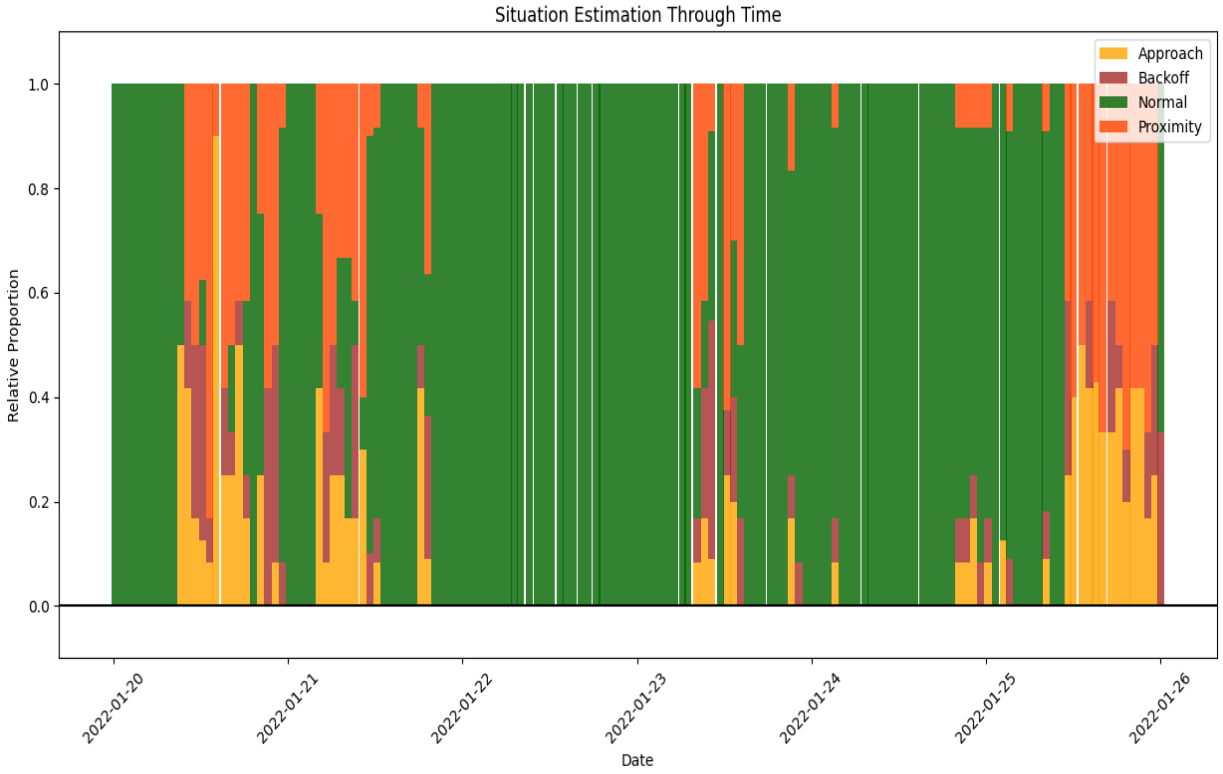

Fig. 5. The fused event timeline giving a situation estimation through time from the E-O processing pipeline output. The y-axis represents the relative proportion of points within the bin sample width (1 hour) that have been classified as a certain event, with different events distinguished by color. For example, around time 2022-01-20 08:00:00 the system reported that half of the reported events were Approach events and half were Normal.

4. Discussion & Future Work

As mentioned in the previous section, there are obvious ways to improve the results including experimenting with and tuning the Extract module's CNN, acquiring or labeling more data for interesting proximity situations to validate the E-O data processing pipeline's performance more explicitly, and using a more sophisticated metric to measure the performance of the Categorize module's multiclassification. This system was intentionally designed to be modular, allowing for easy experimentation in a plug-and-play style.

Other works worth considering incorporating into the system include experimenting with Pruett et al (2023)'s MuyGPs, which enable object discrimination for closely spaced objects [9]. This could be very beneficial if it were employed in the Extract module. Another candidate for eventual integration would be work similar to Vallverdu Cabrera et al (2021), which enabled object shape estimation based on E-O observational data [10].

We've conducted initial experiments with an in-house Focus of Attention algorithm as an alternate to YOLOv5 and are confident there are many other equally performing options that will take future research to incorporate. This also opens the doors to implementing and testing higher level data fusion algorithms, synthesizing multiple event timelines from multiple sensors for the same target for improved SSA.

Finally, an in-depth analysis with regards to common metrics such as the Signal-to-Noise Ratio (SNR) also merit attention. The resulting work would be of a similar approach to Helferty et al's (2023) tracking algorithm performance evaluation [11].

5. Conclusion

In this paper, we presented a modular pipeline for Electro-Optical (E-O) space surveillance observations that leverages both E-O data and imagery to provide near-real-time event detection and characterization. Our prototype system demonstrates the feasibility of processing E-O data and imagery in tandem to identify interesting proximity events, such as multiple satellites in close proximity. The pipeline consists of multiple neural networks and data fusion algorithms to extract, detect, and categorize scenarios from E-O data and imagery. While the current iteration of the prototype still has room for improvement, it has shown promising results in detecting and characterizing threat events.

Future work includes refining the object detection module, increasing the precision and recall of the categorization module, and incorporating additional data sources and algorithms to improve the overall performance of the system. We believe that this approach has the potential to contribute significantly to Space Domain Awareness and enhance the ability of space operators to respond to emerging threats.

6. Acknowledgements

The authors acknowledge funding of this development from the Air Force Research Laboratory (AFRL) under Contract FA8750-23-C-0107. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of AFRL or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

[1] J. Fletcher, I. McQuaid, P. Thomas, "Feature-Based Satellite Detection using Convolutional Neural Networks," in 2019 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[2] J. Zachary Gazak, Ryan Swindle, Matthew Phelps, Justin Fletcher, "SimultanE-Ous Detection, Recognition, and Localization of GE-Osynchronous Satellites from Ground Based Imagery," in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[3] Jacob Lucas, Trent Kyono, Michael Werth, Justin Fletcher, Ian McQuaid, “Automated Resolution Scoring of Ground-Based LE-O Observations Using Convolutional Neural Networks,” in 2019 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[4] Douglas Woodward, Celeste Manughian-Peter, Timothy Smith, Elizabeth Davison, “Pixelwise Image Segmentation with Convolutional Neural Networks for Detection of Resident Space Objects,” in 2021 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[5] Calum Meredith, Paul Chote, Robert Airey, “Tracking Merged Objects within Non-Resolved Imagery,” in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[6] LE-Onid Shakun, Krzysztof Kamiński, Oleksandr Briukhovetskyi, Justyna Gołębiewska, Edwin Wnuk, Monika K. Kamińska, Mikołaj Krużyński, Vadym Savanevych, Oleksandr Kozhukhov, Vladimir Vlasenko, Eugen Dikov, Artem Dmytrenko, Nikolay Koshkin, “Comparison of the LEMUR and PSST Image Processing Pipelines for Astrometric Measurements of Resident Space Objects in All Orbital Regimes,” in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[7] Michael Ian Stewart, Dr. Regina Lee, Shane Ryall, “Image Processing Techniques for Space Situational Awareness,” in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[8] Linesh Patil, Siddhi Khanvilkar, Ashish Shethia, Tulika Jain, Vidyulatta Devmane, Srikanth Kodeboyina, “Threats Prediction to a Satellite by Detected Asteroids,” in 2021 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[9] Kerianne Pruett, Nathan McNaughton, Michael Schneider, “Closely-Spaced Object Classification Using MuyGPyS,” in in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[10] David Vallverdu Cabrera, Jens Utzmann, Roger Forstner, “Inversion of the shape of space debris from non-resolved optical measurements within SPOOK,” in 2021 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).

[11] J. Helferty, J. Smith, E. Ryan, and M. Thomas, “Performance Analysis of Satellite Tracking Algorithms in Low SNR Environments," in 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS).