Space Object Identification and Change Detection Methods for the Cislunar Orbit Regime

Jeffrey Hollon, Kimberly Kinateder, Victoria Carone, Tamara Payne, Markus Ernst, Phan Dao
Applied Optimization Inc.
Charles J. Wetterer
Pacific Defense Solutions (KBR Inc.)
Stephen Gregory
Stephen A. Gregory LLC.
Anthony Dentamaro
Anthony Dentamaro LLC.
James Frith, Scott Milster
AFRL/RVSW

ABSTRACT

A critical function of successful Space Domain Awareness (SDA) is to provide robust and confident Space Object Identification (SOI) and Change Detection (CD) capabilities to analysts. We have developed a collection of algorithms that provide these capabilities for space objects in the Low-Earth Orbit (LEO) and Geosynchronous Earth Orbit (GEO) regimes when surveyed by ground-based and space-based sensor platforms. The proliferation of space objects outside of GEO demands that existing and new technologies address cislunar SDA. Due to differences between the LEO, GEO, and cislunar space domains, existing technologies cannot be applied directly and must evolve to operate in this new and highly dynamic environment. In this paper, we describe the operation of our current CD methods, the problems they face in generalizing from LEO and GEO to cislunar space, and our initial efforts and assessments to extend the CD methods to cislunar space.

The photometry-based CD methods we have developed can be differentiated by how much historical data they require. The baseline method assumes previous photometry data collected on the satellite are available for processing, while the baseline-less method requires only one collection of photometry data. These two methods are complementary in the sense that if the requirements for the baseline method are not met, changes can be assessed using the baseline-less method.

The baseline method ingests historical data on an object that can be from multiple sensors and embeds the astrometric and photometric information in a high-dimensional manifold that statistically learns the space object's brightness characteristics as illuminated and viewed from any direction and with a specific attitude profile (e.g., Earth-nadir pointing). Thus, our approach uses a fusion of astrometric and photometric measurements. The manifold is referred to as the space object's historical hypersurface, and it may be queried based on the orbital state vectors of new observations. The queried historical observations form a statistical baseline without the need to observe the space object under the same observing conditions. These extracted baseline light curves are statistically modelled to provide an expected light curve against which the new data will be evaluated in order to invoke our CD methods.

The baseline-less method comprises a set of three algorithms designed to detect near real-time changes. They differ in that they utilize local trends in the photometric data from one collection, in order to develop an understanding of the space object's behavior. When the local trends are broken, based on statistically driven thresholds, the baseline-less method will flag that a change has occurred.

To demonstrate and test our two CD approaches on objects in cislunar space, we generated a set of notional ground- and space-based photometric observations of targets in selected cislunar orbits. Using the Satellite Visualization and Signature Tool (SVST), we then simulated their photometric light curves. Changes in the target's pose are imposed in the simulations to mimic real-world attitude control changes that cislunar space objects are expected to perform. The baseline and baseline-less methods were tested on this dataset, and their ability to correctly detect these changes in space object behavior was quantified to demonstrate their level of success. Further, we demonstrated that the use of our baseline method provides the ability to perform CD along with SOI.

1. INTRODUCTION

There are many methods and tools available to analyze SOs in the common LEO, GEO, and similar orbit regimes. The relative closeness of these targets to Earth makes acquiring sensor data—optical, radar, and other modalities—easy to accomplish with reasonable accuracies and precisions. Various CD and SOI methods have been developed and tested for these nearby sensors and objects and proven effective in operational environments. The effectiveness of such tools is also aided by the simple set of well-studied Keplerian orbits that these objects assume. Much of our prior knowledge and simplifying assumptions break down when we venture into analyses of cislunar objects.

As scientists we are always looking for new, innovative, and better ways to perform tasks. We are also interested in minimizing the time and effort to make improvements as time and money are both critical aspects of SDA in assuring tactical superiority. As the proliferation and exploitation of the cislunar regime increases, we are faced with several concerning issues for our existing CD and SOI algorithms. The lack of sensors in cislunar space and great distance from Earth makes accurate and precise observation data collection an issue. Furthermore, we find ourselves only at the initial stages of getting space-based sensors online in the cislunar region. Thus, we expect that high-quality data will be sparse and difficult to obtain in the near future. Even if high-quality observations can be obtained, the highly dynamic multi-body physics involved in cislunar orbits makes tracking these targets immensely more complicated. Today we are fortunate to only have a handful of targets that must be tracked and analyzed. However, as growth continues, we foresee a time in the not too distance future where cislunar space may become crowded with active satellites and debris. It is critical for scientific efforts to focus on upgrading existing technologies to enable efficient and effective cislunar SDA so that we do not fall behind in the new cislunar space race. To accomplish this, we have already begun to look at our own technologies and how to transfer their algorithms and knowledge from the LEO and GEO realm into cislunar space. In this paper, we examine two fundamentally different ways of achieving CD and SOI in the LEO and GEO realms that can be extended for exploiting cislunar data. We demonstrate the technology using simulated datasets.

2. BACKGROUND

Optical measurements and photometry play an important role in identifying, maintaining custody, characterizing and monitoring the status of cislunar objects. In this paper, the cislunar regime is interpreted as the space around the Moon and the Earth-Moon L1 and L2 Lagrange points. As optical signatures depend on the Sun as the source of illumination, their availability cannot be readily controlled by the spacecraft operators the way passive radio frequency signatures can. Also, spacecraft attitude changes directly modulate the object's light curve with any changes of status related to the spacecraft's attitude likely being observed in its light curve. Changes in the light curve are what we seek to detect with our algorithms.

The ability to optically detect change in light curves depends on the persistence of the signature collection. Longer and higher cadence light curves are more useful for CD than shorter and sparser windows of observation. The periodicity of the cislunar light curve, generally driven by the lunar day, offers a potentially large window of observation if not disrupted by the various interferences due to bright sources or blockages of the line of sight. For cislunar targets, ground-based optical observations can be interrupted by the daytime hours, weather outages, moon glow, and unfavorable solar phase angles. The daytime gaps can be remediated by using a collaborative network of ground sensors with longitudinal diversity or by deploying the sensors to space. The unpredictable weather gaps can be addressed by having back up sensors. The high background noise zone around the Moon or unfavorable solar illumination can affect the signal-to-noise ratio and introduce observation gaps.

The authors have been developing SOI and CD tools for operational SDA for more than a decade. Until recently, the majority of this work has been in the relatively predictable LEO and GEO regimes. Now we wish to evolve some of our existing technologies and understanding so that they may provide successful SOI and CD tools for targets in cislunar space. Two methods are described in detail within: baseline-less and baseline CD. While both concepts achieve the same end goal (CD), they do so in vastly different approaches. The baseline-less approach builds up very short periods of trend information in a dataset and then attempts to predict the next observation. If the next observation matches sufficiently well with the predicted value, then a change is not likely. The baseline approach examines large sets of data on a particular target and *learns* its appearance as a function of several descriptor angles relating to how the target has been previously illuminated and observed. This learning occurs in a high-dimensional manifold we call the hypersurface. This hypersurface acts as a database that can be queried to obtain expected observations to compare against new, incoming observations and is successful in making meaningful comparisons even for incoming data from observation conditions never previously seen.

For our purposes, we are focusing on electro-optical data as the observations collected on a particular target. These observations are plotted as a time series that we refer to as a light curve. LEO and GEO targets have typical lengths associated with their time series. For example, LEOs make overhead passes on the order of minutes, while GEOs have a natural approximate 12-hour observation window. In the case of cislunar targets, however, there is no natural one-size-fits-all length for such a light curve as the orbital complexities are so numerous. Our simulated light curves are chosen for their illustrative purposes and not based on any real cislunar data. These simulations are done utilizing our experience with orbit generation with NASA's General Mission Analysis Tool (GMAT) and electro-optical observation simulations leveraging the Satellite Visualization and Signature Tool (SVST).

3. PHOTOMETRIC ANALYSIS OF CISLUNAR LIGHT CURVES

In previous work by [1], they presented sample light curves for a diffuse sphere in each of the thirty-one cislunar periodic families and came to the following conclusions for Earth-based sensors: the target's brightness approximately matches the solar phase angle of the Moon: i.e. full/new moon corresponds to peak/trough brightness, global extrema differ from cycle to cycle and correspond to changes in season, and objects in cislunar orbits are fainter by approximately 5 magnitudes as compared to similar objects in GEO due to increased distance. As such, brightness changes in a cislunar object's light curve due to its orbital motion occur over longer time scales that LEO and GEO (days vs. hours) and more rapid changes would be the result of glinting or attitude changes for more complex shaped objects.

3.1 VISIBILITY ANALYSIS

In this paper, we have simulated various Halo orbits and used more complex satellite models to simulate the visibility of a complex object from a space-based sensor. We have also used these simulations to investigate how to extract the object's albedo-Area with photometry from highly diverse illuminations found in the cislunar regime. We used three different radiometric simulations as described below.

The simulation geometries chosen for this work included sensors placed in a variety of cislunar obits and targets in either Halo orbits or translunar/trans Earth trajectories. In order to get statistically significant results, many combinations of observer/target orbits were generated. Also, for each orbit, 10 phases were chosen for each sensor and target orbit, resulting in many millions of simulation steps. An initial simulation plan involved targets and sensors both placed in a variety of 25 Halo orbits, including Halo Northern and Southern orbits about L1 and Halo Northern orbits about L2. In addition, sensors were also placed in two Lyapunov orbits. Given the high-fidelity nature of SVST and its satellite models, executing this full ensemble of scenarios would suggest an exorbitant amount of computer time. So, low-fidelity platforms, on which the entire simulation plan could be run in a reasonable amount of time, were employed as a guide to identify those sensor/target orbit combinations which were most interesting and practical for the initial SVST runs.

Two platforms were used to generate the many thousands of light curves needed for this work. Forge is a facet-based package that includes shadowing and obscuration, and the Analytical Integrated Model Surfaces (AIMS) platform is an analytical code that simply identifies and quantifies the contribution from each model surface. Because AIMS samples entire surfaces instead of individual sub-facets and ignores obscuration, it runs very quickly. All simulations were run on both platforms, and the results were combined to determine which simulations SVST should approach first. As an example, Fig. 1 below compares the visibility of an ANGELS-type satellite, averaged over all target orbits and all sensor/target phases, as predicted by each simulation platform. The satellite was placed in an Earth nadir-aligned pose, and a sensor limiting magnitude of $m \leq 19$ was assumed for the sensor. Although the magnitudes of visibility fractions vary between platforms, the trends predicted by Forge and AIMS were seen to reproduced by SVST and were therefore useful in expediting the generation of results.

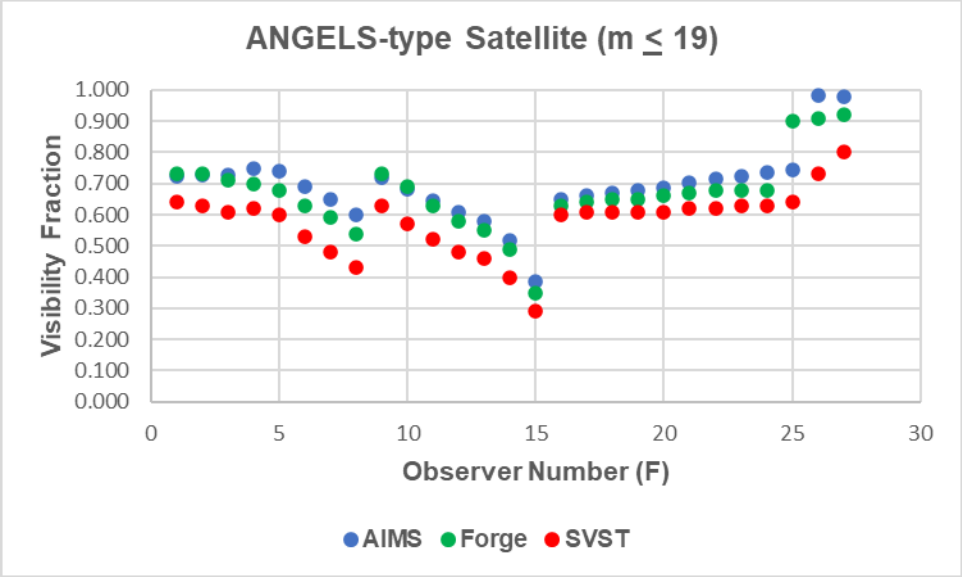

Fig. 1. Averaged Visibility Fractions for an Earth-facing ANGELS-type Target in Various Halo Orbits, Observed by Sensors in Cislunar Orbits, as Predicted by Forge, AIMS and SVST

3.2 CISLUNAR ALBEDO-AREAS

Using the set of simulated magnitudes, we investigated an algorithm for estimating the albedo-Area (aA) product for Resident Space Objects (RSOs) in cislunar orbit. The development is based on models and simulations using Johnson R band magnitudes and ECI (x, y, and z components) as generated by AIMS. We tested a k-Nearest Neighbors (kNN) algorithm in which we adopt a given simulation as a "catalog" and compare the magnitudes and spatial coordinates to those of a second simulation where the second simulation is treated as the "test" object.

In its current form, the algorithm takes each test pair of X and Y Phase Angles (PA) in those specific ECI coordinates with respect to solar illumination associated with each brightness value (in magnitudes). The Pas are defined such that the "full" phase has a value of 0^o coordinates. The algorithm searches the catalog data for the k = 15 nearest neighbors. Each magnitude is normalized to an absolute magnitude using 384,400 km as the standard distance. This catalog absolute magnitude is then converted into a relative flux value, and the mean of the "k" fluxes is evaluated. Next, a ratio of the values of $aA_{test}/aA_{catalog}$ is formed and, for display purposes only, the mean flux is converted back to an absolute R magnitude value.

Fig. 2 presents an illustration of the results for a calculation using an Angels-like model for the catalog and a DIRECTV-1-like model for the test object. In the figure, black symbols represent the test data, blue symbols represent the catalog data, and red symbols represent the KNN mean data as derived from the catalog. We note that this illustration is a specific rotation of a 3-D plot of R_{abs}, ECI_X Phase Angle, and ECI_Y Phase Angle.

In this example, the target's aA is found to be 37 ± 24 times that of the catalog object's aA. When we compare this ratio to that of the "truth" model values of 36.4, we find that this new aA algorithm has a promising level of accuracy at this early stage.

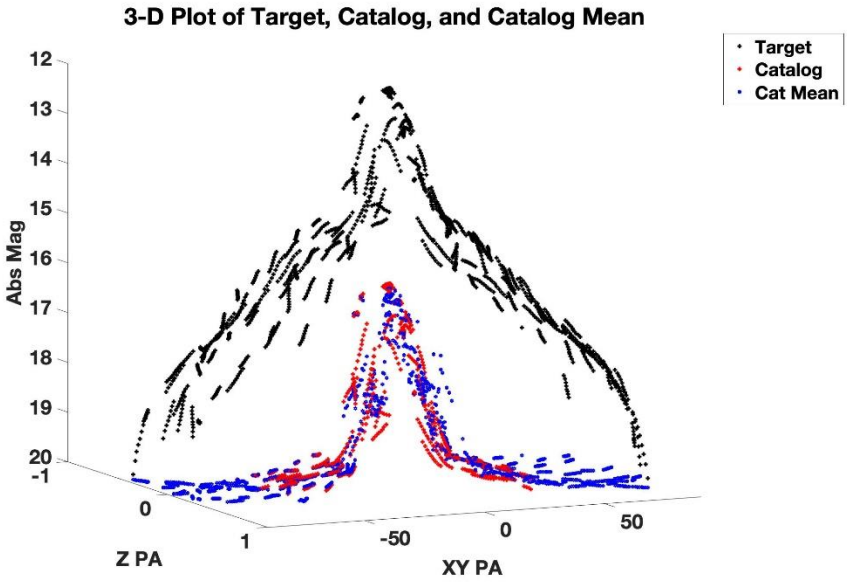

Fig. 2. Comparison of a set of "Catalog" Light Curves (red) and the Catalog Mean Light Curves (blue) with the "Target" Light Curves (black)

4. ALGORITHMS FOR CHARACTERIZATION AND CHANGE DETECTION

Our SOI and CD algorithms take two fundamentally different approaches in how much prior data is needed for them to operate effectively. These two approaches are the baseline and baseline-less methods. The baseline method ingests historical data that has been observed and encodes it into a physics-based manifold called the hypersurface. This encoding allows new incoming observations to be mapped to the hypersurface where nearby historical observations may be extracted to act as an expected mean value for the new data. If the extracted observations are not consistent with the incoming data, then a change is detected. The baseline-less method operates very differently from the baseline method as it only requires limited historical data in order to operate. It assesses the most recent observations in time order and models this local behavior with simple statistical models. These models are used to predict what the next observation should be and if the difference between what is expected and what occurs is sufficiently large, then a change is detected. These two approaches are discussed in more detail in the following sections.

4.1 BASELINE METHOD

The goal of our baseline method is to utilize historical observations on a target to learn how future light curves are likely to appear. Historical observations are ingested and encoded onto a hypersurface in such a way that observations are grouped by observation and illumination similarities. When a new set of observations is observed, the object's historical hypersurface is queried to find the best matching observations relative to the incoming observation's illumination and observation conditions. These best matching observations are then extracted from the hypersurface and used to model an expected light curve with a statistically derived mean and Prediction Intervals (PIs). These extracted observations are referred to as the baseline for the incoming data. The statistical knowledge gained from modeling the baseline is then compared to the incoming observations in a Sequential Probability Ratio Test (SPRT) to detect sequences of incoming observations that are statistically consistent or inconsistent with the baseline. These SPRT decisions are sent through our End-of-Pass (EoP) test as a final filter to determine if the number of inconsistent observation sequences is statistically larger than what would be expected at random. If so, then the incoming data is assigned a label of ANOMalous (ANOM) relative to the hypersurface. Otherwise, the data is considered NOMinal (NOM). The creation of an object's historical hypersurface requires certain assumptions about its attitude control methods and the complexity of the target's components. We assume that targets will be made up of two primary facets—one involving its solar panels and the other being its body or bus. The attitude control assumptions set the angular framework for where each of these facets are pointing and how they move relative to one another. Different attitude control assumptions can be used to produce different hypersurfaces for the same object and thus provide capability of detecting when the attitude control changes. When comparing

hypersurfaces between different objects, it is possible to perform SOI based on which object's hypersurface produces better baselines for a given dataset. It is important to note that this baseline approach is independent of time as it only concerns itself with similar geometric viewing and illumination conditions and not the time that the observations occurred.

4.1.1 BASELINE HYPERSURFACE CREATION

The fundamental part of this procedure is to create a meaningful historical hypersurface that contains all known observation data on a target organized in such a way that similar observing and illumination conditions are situated near one another, with respect to some well-defined distance. A set of angles for each of the two facets are defined based on attitude control assumptions that describe where the observer and where the Sun are positioned relative to that facet. These positions are described by a satellite-centered coordinate system where an elevation and azimuthal angle from a defined normal vector are used to describe the observer's and illuminator's (Sun's) position relative to the target. These angles are defined for both the solar panel and the body facet, where the positions of the observer and Sun are created by pairs of elevation and azimuth angles. This creates a total of eight angles to describe the relationship between the observer, target, and Sun, for each observation. These angles form a set of descriptors, $\{Oaz_B, Oel_B, Oaz_P, Oel_P, Saz_B, Sel_B, Saz_P, Sel_P\}$, that uniquely describe the physical relationship between the Sun, the object, and the sensor at each observation's epoch. The first letter, "O" or "S", indicates whether the angle describes the observer or Sun's position relative to the facet. Then, either "az" or "el" is used to mark the azimuth and elevation angles. Finally, the subscript "B" or "P" identifies if the angle is with respect to the body or panel facet of the target.

We assume that the two facets are independent of one another and recognize that the position of the Sun has no dependency on the position of the observer. Thus, we represent these eight angles as four pairs of angles grouping them by $\{Oaz_B, Oel_B\}$, $\{Oaz_P, Oel_P\}$, $\{Saz_B, Sel_B\}$, and $\{Saz_P, Sel_P\}$. Each pair provides a unique representation of where the observer (or Sun) is located relative to that particular facet and may be represented as spheres by converting each angle pair into a three-dimensional unit vector. This collection of four spheres creates the geometrical landscape that is our hypersurface for encoding historical observations with similar observing and illumination conditions together.

Each of these conversions places a point on one of the hypersurface's spheres and embeds the observed brightness at that location. One such sphere is depicted in Fig. 3 where each point along the sphere's surface represents a particular observation condition, and the various colors represent different hypothetical observed magnitudes. Such a surface is repeated for all four pairs of angles and the set of four generated spheres is our historical hypersurface.

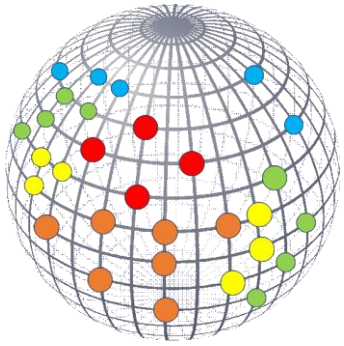

Fig. 3. Illustration of a Single Sphere Encoding Observational Geometry on Its Surface and Colored by Hypothetical Visual Magnitude

4.1.2 BASELINE LIGHT CURVE EXTRACTION AND MODELING

The structure of the hypersurface is used to compare incoming observations to historical ones under similar viewing conditions. To utilize the spheres, the incoming observations are converted to angle pairs and transformed to spherical coordinates that form a path along each sphere contained within the hypersurface. Such a path along one of the spheres is illustrated by the black, dashed curve in Fig. 4.

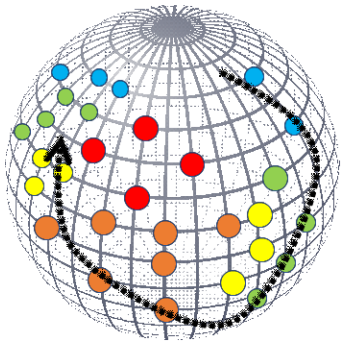

Fig. 4. One Sphere within the Hypersurface with an Incoming Light Curve's Path Plotted along Its Surface

The path that the incoming light curve creates on the hypersphere may pass through nearby historical observations that represent previously observed magnitudes in similar viewing conditions. These historical observations are the ones that are most similar in viewing geometry to that of the incoming light curve and act as a potential baseline of expected magnitudes. For this illustrative example, the k-Nearest Neighbor (k-NN) approach with $k = 1$ was used to select magnitudes to act as the expected values. A possible pair of baseline and incoming light curves is depicted in Fig. 5 where the colored magnitudes are from the hypersphere and the black magnitudes are from the incoming light curve.

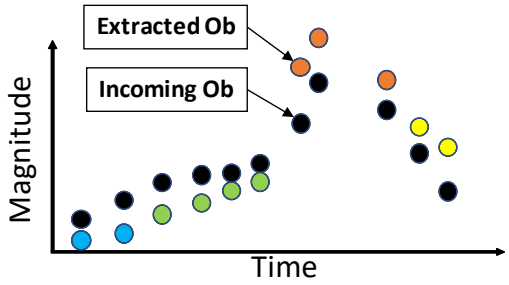

Fig. 5. Sample Extracted Light Curve from the Hypersurface (Colored Circles) along with Hypothetical Incoming Observations (Black Circles)

The use of the k-NN algorithm requires the choice of a distance measure to determine the set of points (observations) to be selected from the hypersurface. Angles between vectors embedded onto the four unit spheres are a natural measure for this manifold, and we use the sum of the four angles across all spheres as our distance metric. The kNN procedure produces a physics-based ordering of historical observations relative to the incoming data and allows for a one-to-k extraction of observations creating a one-dimensional baseline that is directly comparable to the incoming light curve via statistical processing.

4.1.3 STATISTICAL LIGHT CURVE ASSESSMENT

Once hypersurface extraction has completed, the extracted baseline light curve is then statistically modeled to provide quantitative information about its expected behavior. The model fit to the baseline data provides an estimated mean function, uncertainty estimates, and Prediction Intervals (PIs) for the baseline expressing how far from the mean function one could reasonably expect new data to exist within. A statistical learning technique then compares each observation in the new light curve sequentially to the modeled baseline providing a per-sequence assessment of consistency (or inconsistency) between the two observation sets. This sequential analysis returns locations within the light curve where expected changes from the baseline have occurred. The final step in our statistical assessment involves a filtering procedure over all sequential decisions. This final step filters out false positives that can arise from the sequential decision-making process. If more inconsistent regions have been observed than what would be expected from a random experiment, then the entire light curve is flagged as being anomalous with respect to the extracted baseline observations.

The extracted observations that act as expected values for the incoming observations are modeled using a Gaussian Process Regression (GPR) based on [2] and adapted to better represent the expected trends and variability in

photometric data. A careful analytical derivation of the GPR's PIs was done to provide us with confidence as to where randomly sampled observations from the GPR would be expected to appear. This gave us insight into where consistent incoming data would also be expected to appear relative to the baseline light curve. The GPR mean and its PIs are used as the null and alternate hypotheses in a Sequential Probability Ratio Test (SPRT) that we developed based on [3]. This procedure accepts a desired Type I and Type II error rate as input and will sequentially determine if the incoming observations are likely to have originated from the baseline's functional mean (consistent) or the baseline's functional PIs (inconsistent). This sequential procedure is performed for all incoming observations. A novel adaptation of the SPRT procedure that we developed was to add a conditional property to sequential decisions such that the previous decision influences the next one. In doing so, our SPRT procedure is aware of previous inconsistencies and adapts to become increasingly suspicious when multiple inconsistent observations appear in a row. Similarly, the process adapts to consistent sequences of observations, allowing it to relax and not declare changes too often.

The final step in our statistical assessment is a filtering technique designed to determine if the number of inconsistent sequences of observations is greater than what would be expected at random from SPRT. This End-of-Pass (EoP) procedure outputs a p-value for the null hypothesis that the incoming and baseline light curves are consistent with one another. If a sufficient number of inconsistent sequences from SPRT are observed, the p-value will be small, indicating an anomalous situation.

AO's development of the GPR and SPRT technology unlocks a CD method that can achieve confidence when deciding if light curves are either consistent or inconsistent with baseline data using minimal observations. In contrast, standard hypothesis testing procedures can only achieve a confidence in the rejection of consistent observations. Further, our research has provided a method that can filter through sequential decisions made by SPRT and create summarizing conclusions. This is achieved with our EoP test that assigns an overall decision of consistency to an incoming light curve. If EoP indicates a significant change has occurred in a light curve, the SPRT results may be reviewed as a tool to examine the more finely detailed changes in behavior of the incoming dataset.

4.2 BASELINE-LESS METHOD

The goal of our baseline-less method is to build up mathematical knowledge on recent behaviors in the time series of observations to make locally, statistically based predictions for future values. Then, as more observations are ingested, these predictions can be tested. If the predictions turn out to be true, then the new data are labeled as consistent. Otherwise, if the predictions are found to be significantly violated by the incoming observations, then an inconsistency can be flagged, and a likely change detected. There are three algorithms that create our suite of baseline-less CD: the Discontinuity Step (DS), Cumulative Sum control chart (CuSum), and Double Simple Linear Model Regression (DSLMR) algorithms. These algorithms are designed to pick up on different types of changes that are possible in a one-dimensional time series. For example, DS detects abrupt changes in observations, CuSum detects unusual changes in the mean value of the observations, and DSLMR examines changes in the trend of observations through time. While these three algorithms are designed to pick up different changes, they can also pick up similar ones as well. In order to reduce the false positives generated by running three independent algorithms on the same data, a filtering of each detected change is performed utilizing a so-called Interval Of Interest (IOI). An IOI is a specific interval of time where the union of the three baseline-less CD methods have detected a sufficiently dense number of changes. If the density of changes from all methods is below the defined thresholds, then the changes are considered likely to be from random sources and not an actual change in the light curve. All these methods are statistically based algorithms with user-tunable parameters to meet the needs of different SDA applications.

4.2.1 THE THREE BASELINE-LESS ALGORITHMS

Three separate algorithms are incorporated in our baseline-less CD method including the DS, CuSum, and DSLMR techniques. Each of these three are designed to detect different types of behaviors that would be indicative of change in a light curve.

4.2.1.1 DISCONTINUITY STEP

Discontinuity step detects jumps in the observations by examining recent linear trends in the data to see if the next observation matches the expected linear trend. The algorithm proceeds by taking a sequence of observations within a set window and fitting a linear model to them as a function of time. The residuals from this linear model are calculated and used to estimate a threshold for the current observation's expected value. If this threshold is surpassed, then the current point is flagged as a change in the form of a step. DS accepts a Type I error rate and a

window length specifying the length of each observation sequence as inputs. This algorithm and the type of change it is designed to detect are illustrated in Fig. 6.

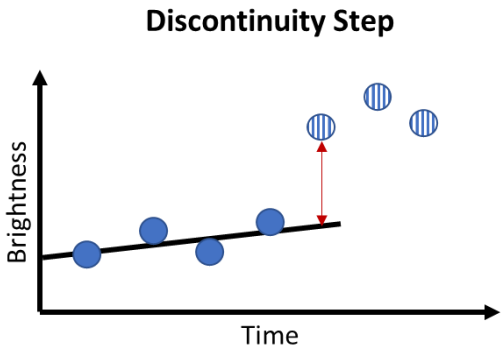

Fig. 6. Illustration of Discontinuity Step Method

4.2.1.2 CUMULATIVE SUM CONTROL CHART

The CuSum method detects a shift in the mean of observations via a control chart method ingesting residuals of recent observations from a low-degree polynomial fit. The method is based on [4] and utilizes a window that slides down the light curve. In each window, a polynomial is fit whose residuals are processed by the CuSum control chart technique. The algorithm accepts Type I and Type II error rates as input along with the window length and the desired amount of change in the mean to detect. Fig. 7 illustrates the CuSum algorithm and the type of change it addresses.

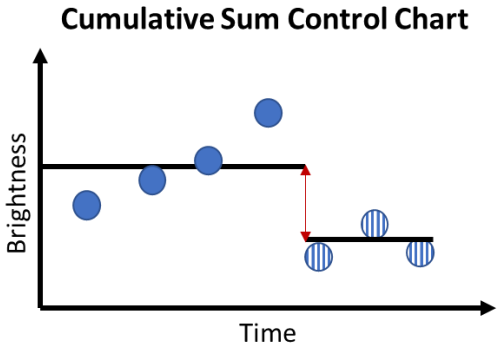

Fig. 7. Illustration of CuSum Method

4.2.1.3 DOUBLE SIMPLE LINEAR MODEL REGRESSION

DSLMR detects a change in the trend of observations by comparing expected observations from both a previous and future small timespan of data. It proceeds down a light curve by selecting a window both before and after the current observation to be analyzed for change. All points in each window are independently fit to a linear model. Predictions for the current observation are estimated separately using the fit from each model. If the difference in the mean estimated value to the actual value is larger than a threshold, then it is flagged as a change in trend. DSLMR accepts a Type I error rate, a window length, and the desired level of change that should be detected as inputs. This algorithm and the changes it is designed to detect are illustrated in Fig. 8.

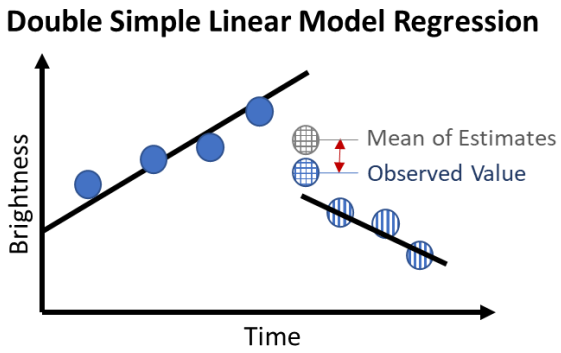

Fig. 8. Illustration of DSLMR Method

4.2.2 THE INTERVAL OF INTEREST

The three CD algorithms are run independently through time on a light curve, and each determine their own change points. In a nominal (NOM) light curve, we expect false positives that are randomly scattered throughout time. In an ANOM light curve, we expect clusters of change points from at least one of the algorithms. The IOI is a way of decreasing the proportion of false positives by only flagging regions of a light curve where a significant number of changes are detected. Restrictions on the minimum number of changes and minimum number of methods that detected a change can be placed on the IOI.

An IOI is defined as any time range satisfying the following conditions:

- The start of the IOI is a change point.
- The end of the IOI is a change point.
- No subinterval of the IOI goes more than 10 minutes without a change point.

Fig. 9 illustrates how the IOIs are determined for a notional light curve when requiring the IOI to contain at least two change points.

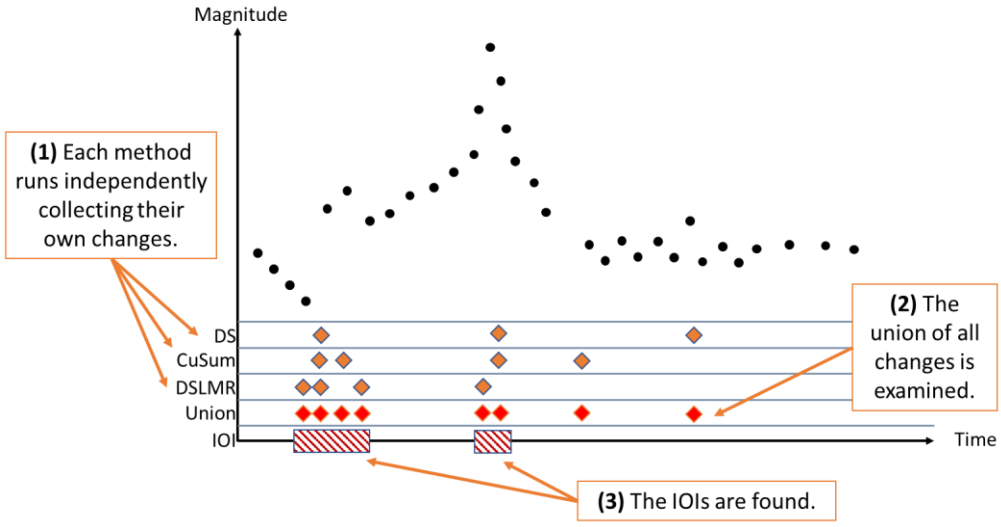

Fig. 9. Illustration of Method to Locate IOIs

5. DATA SIMULATION

GMAT was used to design an inclined lunar orbit with a semi-major axis of 60,250 km. A large semi-major axis was chosen so that ground sensors could view the target outside the Moon's exclusion zone, i.e., cone-of-shame, which was set as 5° from the center of the Moon. The position and velocity information were exported from GMAT and

imported into SVST (version 8.3.61). SVST was used to simulate electro-optical observations—light curves—over the course of four months (August 2019-November 2019) for three different three-axis stabilized targets in the same inclined lunar orbit. These three targets were an Earth-facing Galaxy 14 model, a Moon-facing Galaxy 14 model, and an Earth-facing Inmarsat 4 model. Artist renderings of the Galaxy 14 and Inmarsat 4 satellites are shown in Fig. 10. Two different sensor locations were chosen at geographically diverse locations in order to reduce gaps in the light curve, since only simulated data from when the target was observable were considered. The target was observable if it was nighttime at the sensor's location, the target was not blocked by the Earth or Moon, and the target was not in the Moon's exclusion zone.

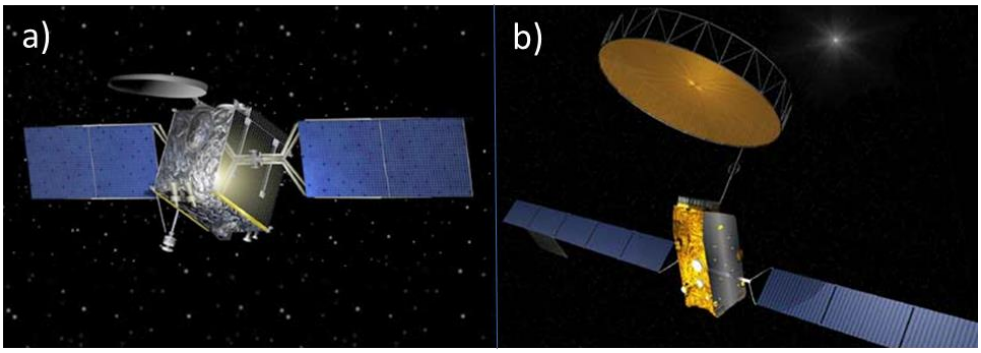

Fig. 10. Artist Renderings of a) Galaxy 14 and b) Inmarsat 4 from Gunter's Space Page[1]

6. RESULTS

The baseline and baseline-less methods were tested on the simulated photometry data. The simulated signatures and the results from the baseline and baseline-less methods are discussed in this section.

6.1 SIMULATED SIGNATURES

SVST was used to simulate light curves over a four-month duration for three different targets in the same lunar orbit. Sensors were place at two different geographically diverse ground-based locations. The three simulation targets include an Earth-facing Galaxy 14 model, a Moon-facing Galaxy 14 model, and an Earth-facing Inmarsat 4 model, and their light curves for the four-month simulation are plotted in Fig. 11. The color bar at the bottom of Fig. 11 represents the lunar phase; the lightest color corresponds to Full Moon, and black corresponds to New Moon. Notice that the behavior of the light curves for the two Earth-facing targets repeats with lunar phase, though the glints are brighter in November. The light curve for the Moon-facing target does not have the same behavior month-to-month, though in general, it is brightest near Full Moon and dimmest around New Moon. The widely varying magnitudes near New Moon are from micro-glints from different components of the satellite models when they are backlit.

[1] https://space.skyrocket.de/index.html

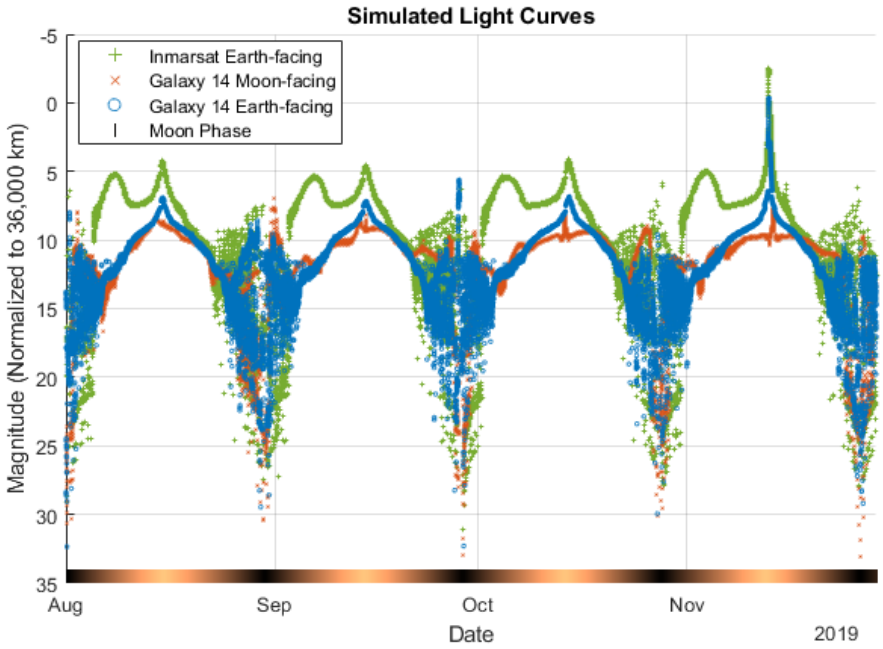

Fig. 11. Light Curves over Four Months for the Three Models with Moon Phase Represented by the Color Bar at the Bottom

Fig. 12 and Fig. 13 show these light curves for smaller time spans: a month and on the order of days, respectively. Fig. 12 shows the light curves for the month of October. The light curves from the three different models have distinct trends in brightness and different magnitude ranges. All three models produce glints around Full Moon on October 13, 2019, though the shape and maximum brightness depend on the model. These glints are shown in the left panel of Fig. 13; note that the discontinuities seen are due to the data coming from two different sites. There are other portions of the light curves where the brightness values are more similar, for example, on October 18th; these light curves are plotted in the right panel of Fig. 13. While there are more similarities on this day, there are still visual differences between the three light curves. These simulated light curves with distinct trends and brightness ranges illustrate that electro-optical data are exploitable for Change Detection (CD), Space Object Identification (SOI), and Space Object Characterization (SOC).

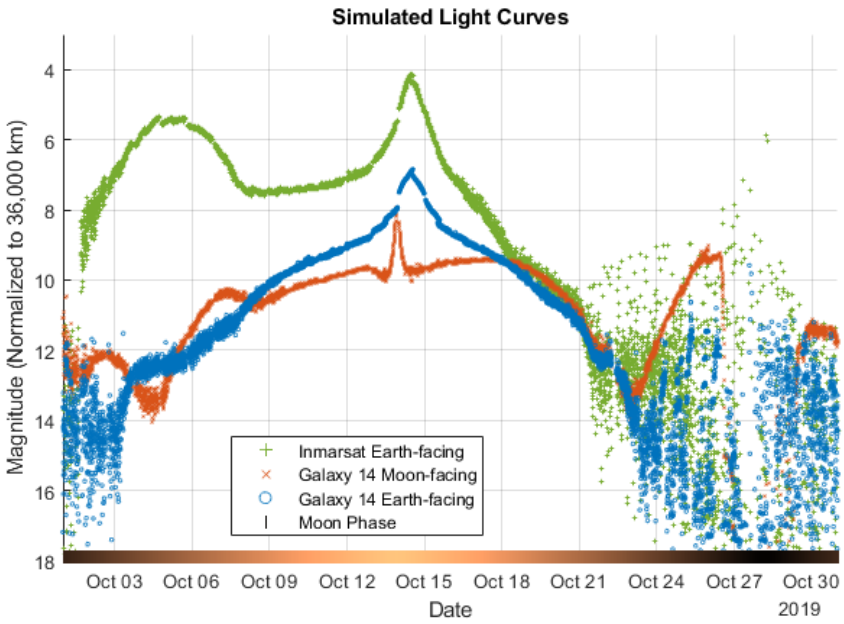

Fig. 12. October Light Curves for the Earth-facing Galaxy 14, Moon-facing Galaxy 14, and Earth-facing Inmarsat 4 Models

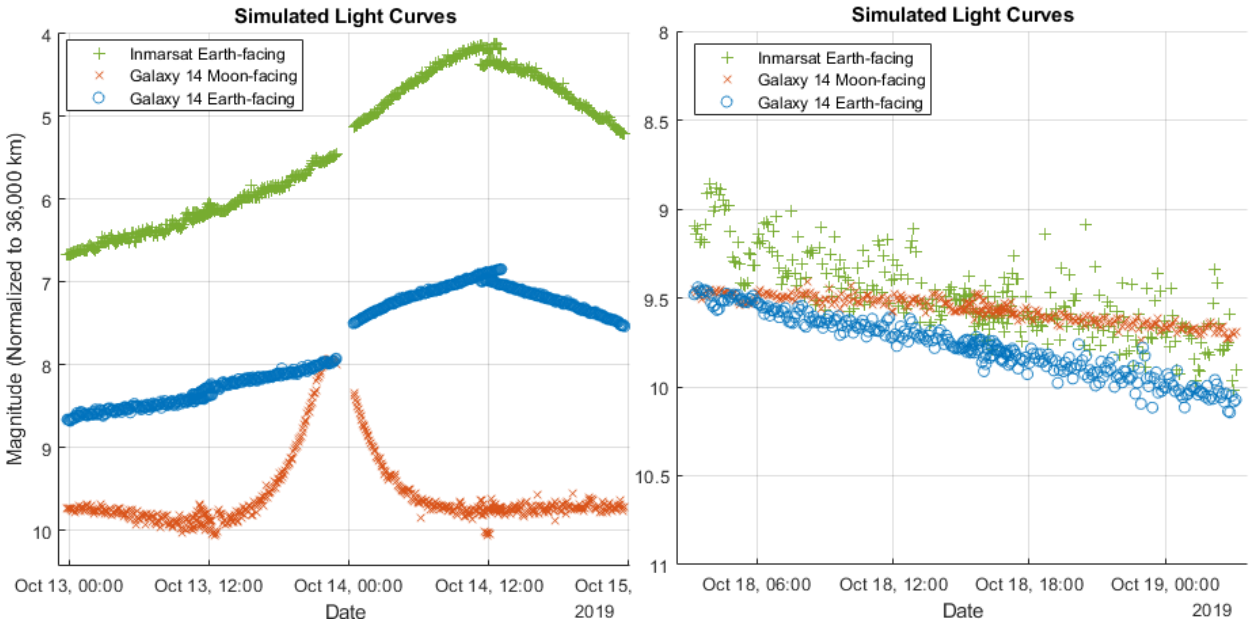

Fig. 13. Light Curves for October 13-14th (Left) and October 18th (Right) for the Earth-facing Galaxy 14, Moon-facing Galaxy 14, and Earth-facing Inmarsat 4 Models

6.2 BASELINE METHOD

This section includes preliminary results from the baseline method for simulated nominal (NOM) and anomalous (ANOM) light curves from October 18, 2019. There are two steps to the baseline method: (1) extract baseline data from the hypersurface and (2) use Gaussian Process Regression (GPR) + Sequential Probability Ratio Test (SPRT) to determine if the incoming optical data are NOM or ANOM. Note that the hypersurface allows for data from two different sensors to be combined, which is extremely useful for ground-based sensors that have a limited field of regard and can only collect optical data at night. Fig. 14 shows an Earth-facing Galaxy 14 incoming light curve (black dots), and the baseline (colored x's) extracted from the hypersurface built using historical Earth-facing

Galaxy 14 data. The color of the baseline observations represents the distance between the incoming and baseline observations as measured on the hypersurface. The extracted baseline matches the incoming data well as expected for this NOM case. Notice that the baseline observations at the beginning of the light curve that are farther away from the incoming data (colored red) do not match the data as well observations that are closer (colored blue). This matches our intuition that if the observation conditions are more similar, then the brightness values will be more similar.

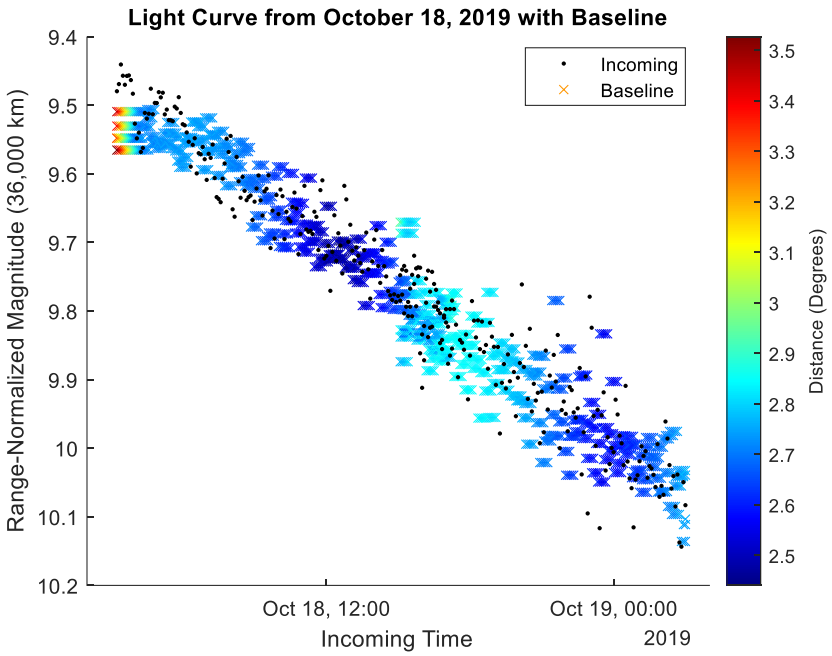

Fig. 14. Incoming Earth-facing Galaxy 14 Light Curve and Extracted Baseline from Earth-facing Galaxy 14 Hypersurface

The baseline light curve was fit using our custom GPR model, and the associated Prediction Intervals (PIs) for the GPR fit were calculated. As calculated, these PIs should include 99.999% of the extracted baseline observations. Our SPRT procedure utilized these PIs to make a comparison with the incoming observations and determine if the incoming data were NOM or ANOM; a decision was only made if there was sufficient evidence to make it at a set confidence level. The incoming light curve colored based on the SPRT results and the GPR fit and PIs for the baseline are plotted in Fig. 15. All decisions were NOM except one, and the End-Of-Pass (EOP) test returned a statistical p-value of 0.9858. The p-value represents the consistency of the incoming data to the baseline data, where higher values, with a maximum of one, indicate the data are more consistent. Thus, this p-value of 0.9858 suggests the incoming light curve matched the baseline, and the baseline method successfully determined the incoming data were NOM.

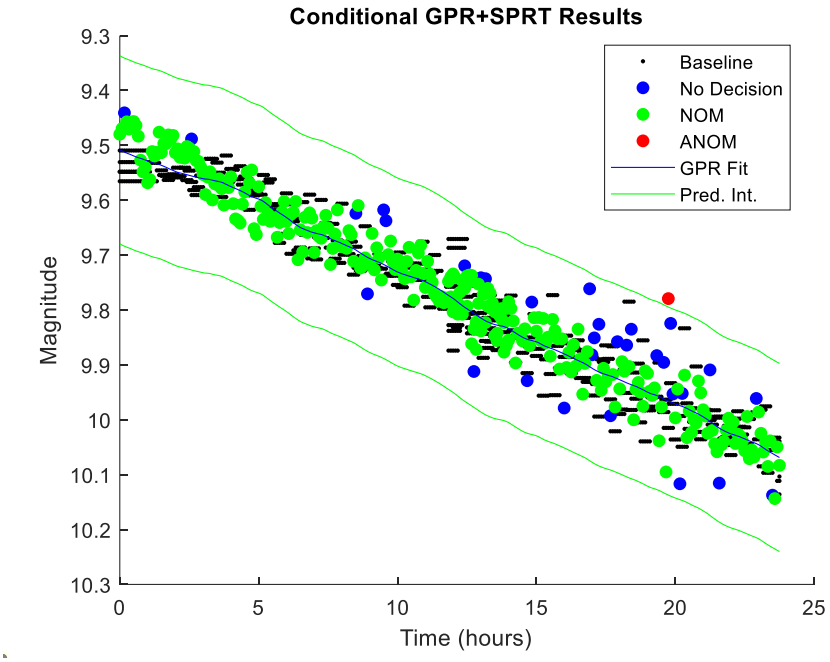

Fig. 15. GPR+SPRT Results for Earth-facing Galaxy 14 Light Curve Compared to Baseline from Earth-facing Galaxy 14 Hypersurface

Next, the baseline method was used on an ANOM light curve where the Earth-facing Galaxy 14 target switched to Moon-facing partway through the night after 8 hours and 40 minutes. This ANOM incoming light curve was compared to baseline data extracted from the hypersurface built using historical Earth-facing Galaxy 14 data. Thus, the incoming data should be considered NOM at the beginning of the night, then after the attitude change, the incoming data should be ANOM. Fig. 16 shows the GPR fit and PIs for the baseline data and the incoming light curve colored based on the SPRT decisions. Before the attitude change, all the SPRT decisions were NOM; however, after the attitude change, all the SPRT decisions were ANOM. The EOP p-value was 7.3528e-11, indicating a strong rejection that the signature is NOM. Thus, the baseline method correctly determined the incoming light curve was ANOM and successfully flagged when the ANOM behavior began.

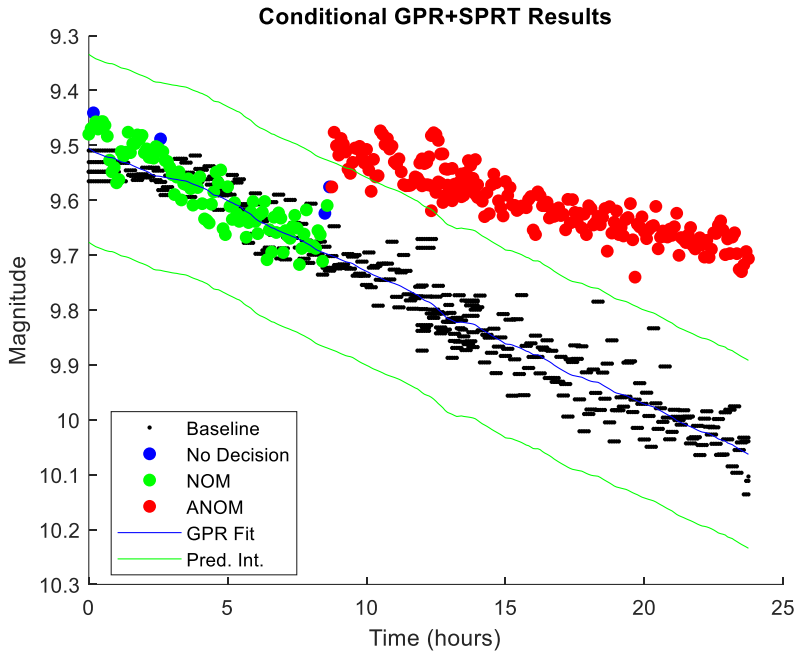

Fig. 16. GPR+SPRT Results for Earth-to-Moon-facing Galaxy 14 Light Curve Compared to Baseline from Earth-facing Galaxy 14 Hypersurface

The baseline method was then used for a cross-tag scenario where two objects are not properly deconflicted. In this case, the incoming observations were initially from the Earth-facing Galaxy 14 target, then partway through the night, the observations were changed to originate from the Inmarsat 4 target. The incoming signature was tested against baselines extracted from two different hypersurfaces: Earth-facing Galaxy 14 and Earth-facing Inmarsat 4. Fig. 17 shows results for the baseline from the Earth-facing Galaxy 14 hypersurface on the left and from the Earth-facing Inmarsat 4 hypersurface on the right. GPR+SPRT correctly detected the ANOM behavior in both cases. On the left of Fig. 17, the Inmarsat data are determined to be ANOM, and on the right, the Galaxy 14 data were flagged ANOM. The EOP p-values were less than 10^{-3} in both cases, indicating the incoming data were ANOM. Note that while the majority of the ANOM portion of the incoming light curve was within the PIs on the right in Fig. 17, the data were consistently on one side of the GPR mean and toward the edges of the PIs, so SPRT was able to detect the ANOM behavior.

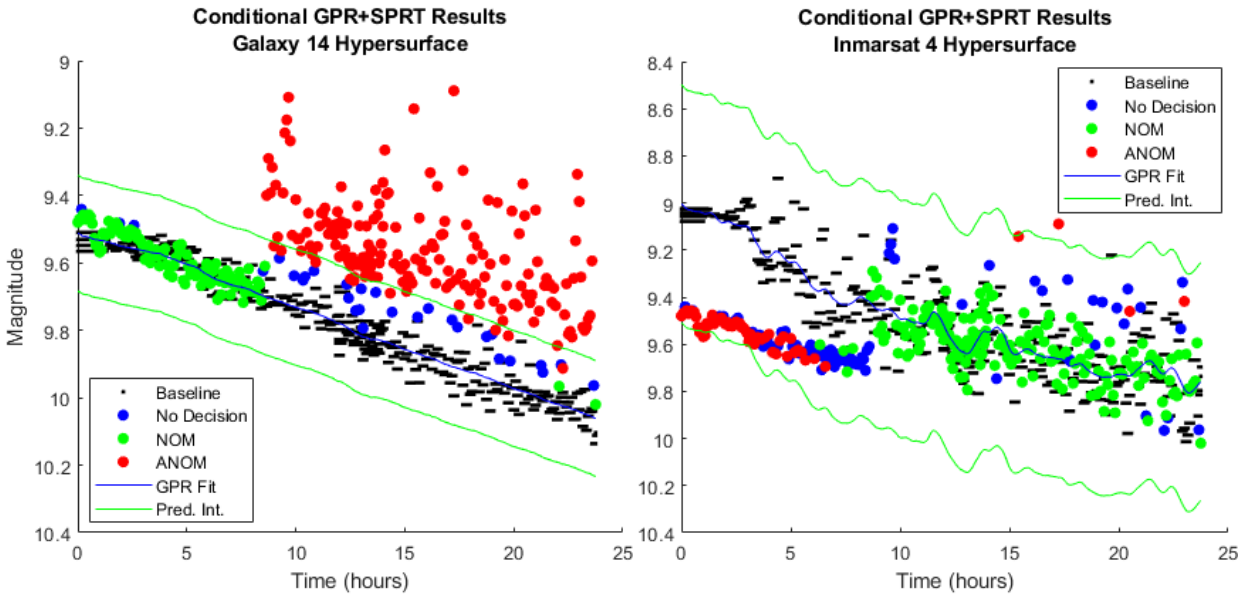

Fig. 17. GPR+SPRT Results for Earth-facing Galaxy 14 Then Inmarsat 4 Cross-tag Scenario Compared to Baselines from Earth-facing Galaxy 14 Hypersurface (Left) and Earth-facing Inmarsat 4 Hypersurface (Right)

The baseline method was able to detect ANOM behavior, but can it be used to determine which object's data are cross-tagged? To test this, the incoming observations from the Earth-facing Inmarsat 4 target after the cross-tag were compared to baseline data extracted from the hypersurfaces built using historical Earth-facing Galaxy 14 data and Earth-facing Inmarsat 4 data. The GPR+SPRT results for this test are plotted in Fig. 18. The incoming Inmarsat 4 data match the extracted Inmarsat 4 historical observations much better than the historical data of Galaxy 14. The EOP p-value for the Galaxy 14 hypersurface was 6.1154e-11, and the p-value for the Inmarsat 4 hypersurface was 0.8515. From these results, we correctly conclude the cross-tagged data were from the Inmarsat 4 target.

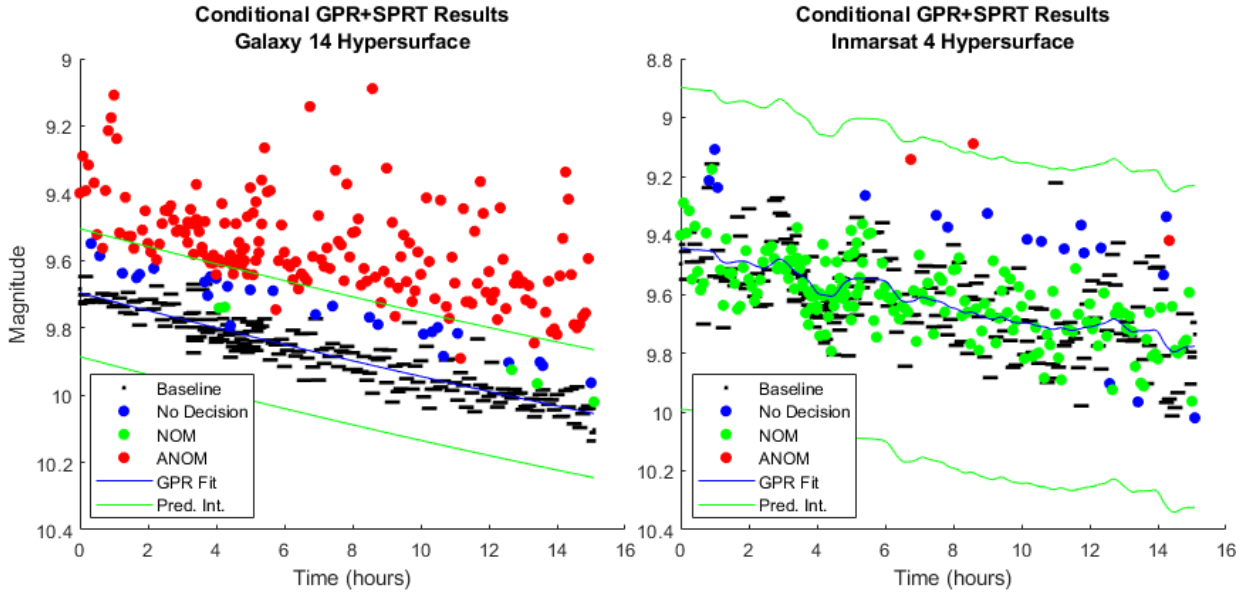

Fig. 18. GPR+SPRT Results for Earth-facing Inmarsat 4 Light Curve after Cross-tag Compared to Baseline from Earth-facing Galaxy 14 Hypersurface (Left) and Earth-facing Inmarsat 4 Hypersurface (Right)

6.3 BASELINE-LESS METHOD

This section presents preliminary results from the baseline-less method. The three baseline-less algorithms work near-real-time, so the incoming data are evaluated sequentially. If at least two incoming observations within a certain window are flagged as ANOM and these changes result from at least two of the three baseline-less algorithms, the observations are labeled an IOI.

First, the light curve and baseline-less results for the NOM Earth-facing Galaxy 14 target are plotted in Fig. 19. Only observations from one site were used. A few of the observations are falsely identified as changes; however, they all come from the same algorithm, CuSum, so there is no IOI. Thus, the baseline-less method concluded this light curve was NOM as expected. This light curve is a subset of the light curve shown in Fig. 15, and the baseline-less method's results are consistent with the baseline method's results.

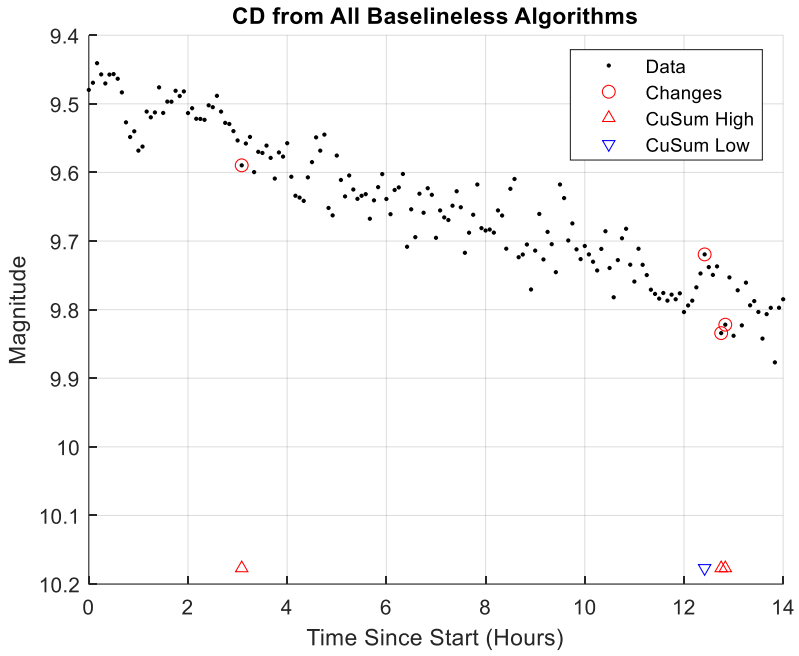

Fig. 19. Baseline-less Method Results for Earth-facing Galaxy 14 Light Curve

The baseline-less method was then tested on an ANOM case where the Galaxy 14 target's attitude changed from Earth-facing to Moon-facing after 8 hours and 40 minutes. The light curve and baseline-less results are plotted in Fig. 20. Two of the three baseline-less algorithms detected the attitude change, and the attitude change was included in the IOI. Thus, the baseline-less method correctly identified the ANOM behavior. This is consistent with the baseline method's results shown in Fig. 16.

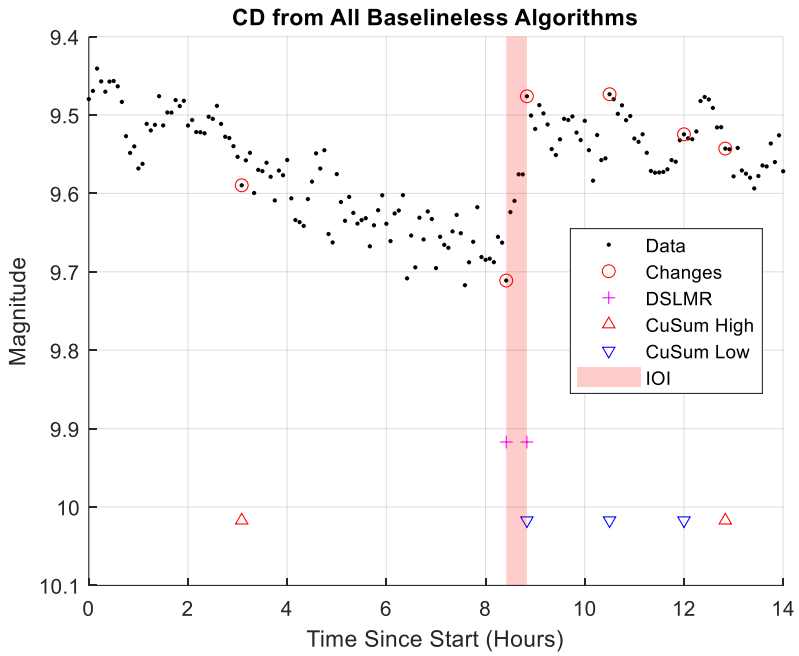

Fig. 20. Baseline-less Method Results for the Earth-to-Moon-facing Galaxy 14 Light Curve

The baseline-less method was next used for the cross-tag scenario where the observations before 8 hours and 40 minutes are for the Earth-facing Galaxy 14 target and the data after are for the Earth-facing Inmarsat 4 target. The baseline-less results are plotted in Fig. 21. The time of the cross-tag is correctly included in an IOI that contains change detections from all three algorithms. While the baseline-less method successfully detected the change, it cannot determine the origin of the ANOM data. The baseline method is necessary for identifying the cross-tagged object.

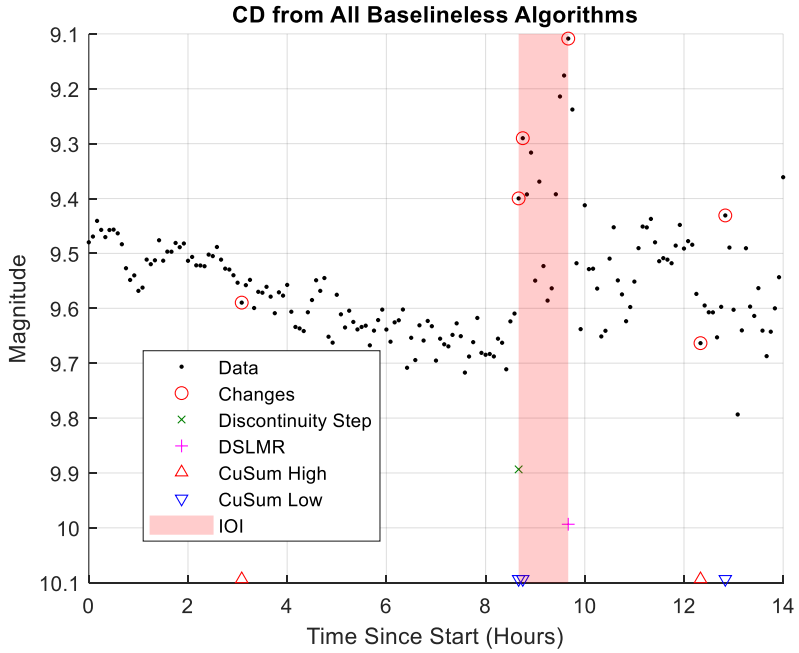

Fig. 21. Baseline-less Method Results for Earth-facing Galaxy 14 then Earth-facing Inmarsat 4 Light Curve

In summary, these preliminary results demonstrate that both the baseline and baseline-less methods are effective at detecting changes in XGEO targets' behavior. If there are sufficient historical data available, the baseline method

can detect changes as well as identify and characterize an XGEO target. As mentioned in the introduction, we expect limited availability of high-quality data (data with high photometric precision) will be sparse and difficult to obtain on cislunar targets in the near future. Thus, in cases of no historical data or no historical data for sufficiently similar conditions, the baseline-less method can be used to detect changes. However, with no pattern-of-life information, sudden changes in brightness may be falsely flagged as ANOM by the baseline-less method.

7. CONCLUSIONS

The authors have completed simulations to better understand how observing and illumination conditions affect electro-optical data trends in the cislunar regime. These simulations demonstrated that light curves from different satellites and attitude modes have distinct trends in brightness and different magnitude ranges allowing for Change Detection (CD), Space Object Identification (SOI), and Space Object Characterization (SOC). These simulations provided a basis for our existing CD algorithms to be tested once evolved to operate in the cislunar regime. The results presented provide proof-of-concept that our baseline and baseline-less CD methods, originally designed for LEO and GEO analysis, can perform CD, SOI and SOC on cislunar targets.

Two different methods have been developed and modified to perform SDA in the cislunar regime using photometry data. The methods can be differentiated by how much historical data they require. The baseline method assumes previous data collected on a satellite are available for processing, while the baseline-less method requires only one collection of photometry data.

The baseline method encodes historical data on a physics-based hypersurface to determine what historical data are from sufficiently similar observation and illumination conditions. A baseline is extracted from the hypersurface and compared to incoming data using our custom GPR model and conditional SPRT process. AO's development of a custom GPR and SPRT technology unlocks a CD method that can achieve confidence when deciding if light curves are either consistent or inconsistent with baseline data using minimal observations. In contrast, standard hypothesis testing procedures can only achieve a confidence in the rejection of consistent observations. Further, our research has provided a method that can filter through sequential decisions made by SPRT and create summarizing conclusions. This is achieved with our EoP test that assigns an overall decision of consistency to an incoming light curve. If EoP indicates a significant change has occurred in a light curve, the SPRT results may be reviewed as a tool to examine the more finely detailed changes in behavior of the incoming dataset.

The baseline-less method is a collection of three algorithms designed to detect different types of changes in local temporal trends. The three baseline-less CD algorithms are run independently through time on a light curve, and each determine their own change points. The IOI was defined so that only sections of a light curve where that are a significant number of changes are flagged as a region of interest. The baseline and baseline-less methods are complementary in the sense that if the requirements for the baseline method are not met, changes can be assessed using the baseline-less method.

The results presented within provide a proof-of-concept that the baseline and baseline-less methods can successfully detect changes in attitude and cross-tags. This paper also demonstrates that the baseline method can be used for SOI since it provides pattern-of-life information.

8. FUTURE WORK

This paper presented foundational results for future experiments on algorithms for cislunar CD, SOI, and SOC. There are many paths forward that we wish to explore for our demonstrated technologies. The primary desire for future work is to improve the realism of the use cases. For example, the simulation of gradual attitude changes as opposed to abrupt ones would provide knowledge as to how the various methods perform relative to various rates of change in the spacecraft. Another area to improve on realism is the addition of a physics-based noise model to the simulated observation data.

The statistical techniques used within our CD methods were originally designed and optimized for signatures of objects in the LEO and GEO regimes. With the much more dynamic orbital types in cislunar space, we expect more dynamic light curves from these objects on average. These more varied signatures may require fine-tuning to the GPR fitting technique, e.g., in the choice of the smoothing kernel. Other areas for adjustment are the statistical input parameters to the CD algorithms that can be adjusted on a case-to-case basis as needed. Increasing the number of simulations of realistic use cases—or the collection of real example data—will help guide our determination for optimal parameters.

Within the baseline method, a finer analysis of the density of observations embedded on the hypersurface will provide insight into how much data are needed for the method to work adequately. This analysis will also help guide thresholds for what distance metric and associated threshold is ideal for extracting the most meaningful nearest neighbors for baseline creation.

9. ACKNOWLEDGEMENTS

The authors would like to thank the Air Force Research Laboratory (AFRL) for their support of the technology development. We would also like to thank AFRL for SVST (a vital tool in our analysis), and Dr. Gregg Crockett at Tau Technologies for his timely answers to our various questions using SVST for this research.

10. REFERENCES

[1] C. Chow, C. Wetterer, K. Hill, C. Gilbert, D. Buehler and J. Frith, "Cislunar Periodic Orbit Families and Expected Observational Features," in *AMOS Technologies Conference*, Wailea, 2020.

[2] C. E. Rasmussen and C. K. Williams, Gaussian Processes for Machine Learning, Cambridge, MA: MIT Press, 2006.

[3] A. Wald, Sequential Analysis, New York: J. Wiley & Sons, 1947.

[4] D. C. Montgomery, Introduction to Statistical Quality Control, Hoboken, NJ: John Wiley & Sons, Inc., 2009.