Geosynchronous satellite maneuver classification via supervised machine learning

Thomas G. Roberts
Massachusetts Institute of Technology
Richard Linares
Massachusetts Institute of Technology

ABSTRACT

This work describes an approach for detecting the components of longitudinal shift maneuvers in the geosynchronous (GEO) orbital regime using convolutional neural networks trained on publicly available two-line element (TLE) data from the U.S. Space Command's space object catalog. A method for converting TLE data to geographic position histories—longitude, latitude, and altitude positions over time in the Earth-centered, Earth-fixed geographic reference frame—and labeling longitudinal shift maneuvers by inspection is described. A preliminary maneuver detection algorithm is designed, trained, and tested on all GEO satellites in orbit from January 1 to December 31, 2020. Performance metrics are presented for a suite of algorithms trained on data sets corresponding to ten years' worth of geographic position time-histories labeled with longitudinal shift maneuvers. When detected, longitudinal shift maneuvers can be used to identify anomalous behavior in GEO. In this work, a satellite's behavior is considered nominal if it adheres to the satellite's pattern of life (PoL)—its previous on-orbit behavior made up of sequences of both natural and non-natural behavioral modes, including routine station-keeping, other on-orbit maneuvers, and uncontrolled motion—and anomalous if it deviates from the satellite's PoL. Identifying anomalous satellite behavior is of critical interest to space situational awareness system operators, who may choose to task their sensors to obtain more observations of anomalous behavior, and satellite operators themselves, who may wish to diagnose its root cause. Applications of this work for international space policymaking is also discussed.

1. INTRODUCTION

To maintain their position relative to the Earth's surface, satellites in the geosynchronous (GEO) orbital regime must expend onboard propellant to counteract any natural forces that perturb their orbits [38], such as those associated with the non-uniformity of Earth's gravitational field, the Sun and Moon's gravity, and solar radiation [34, 68]. Due to the magnitude of these perturbations, GEO satellites must perform station-keeping maneuvers frequently—on the order of once per week to several times per year for high-thrust chemical propulsion systems—in order to stay near their assigned position [6]. Less frequently, GEO satellites perform maneuvers to alter their orbital characteristics more drastically. One such maneuver—known as a *longitudinal shift*—is associated with changing a GEO satellite's sub-satellite point from one position on the Earth's equator to another. Such a maneuver often requires a series of impulsive thrusts to first remove the satellite from its initial position within the geostationary belt, allow it to naturally drift either eastward or westward around the belt, and finally reposition itself at a new longitudinal position.

An operator may choose to perform a longitudinal shift maneuver for a variety of reasons. A military operator may choose to shift a GEO satellite's sub-satellite point to alter its observable region in response to a changing threat environment on the ground. A commercial communications operator may wish to better accommodate an unexpected change in customers' regional demands. If a GEO satellite is part of a *constellation*—a network of two or more satellites that work cooperatively to meet a set of mission requirements—its operator may choose to reposition it in conjunction with other aspects of the constellation's operational evolution, such as the retirement or addition of another satellite in the network. Lastly, an operator may choose to initiate an eastward or westward drift—the first component of a traditional longitudinal shift maneuver—as a means of retiring a satellite from service, since initiating an eastward or westward drift corresponds to lowering or raising the satellite's orbital altitude, respectively.

Like the most popular orbital regime—low Earth orbit (LEO), in which satellites orbit the Earth at altitudes of 2,000 km or less—the GEO environment has grown more populated over time. Although the launch of more space objects

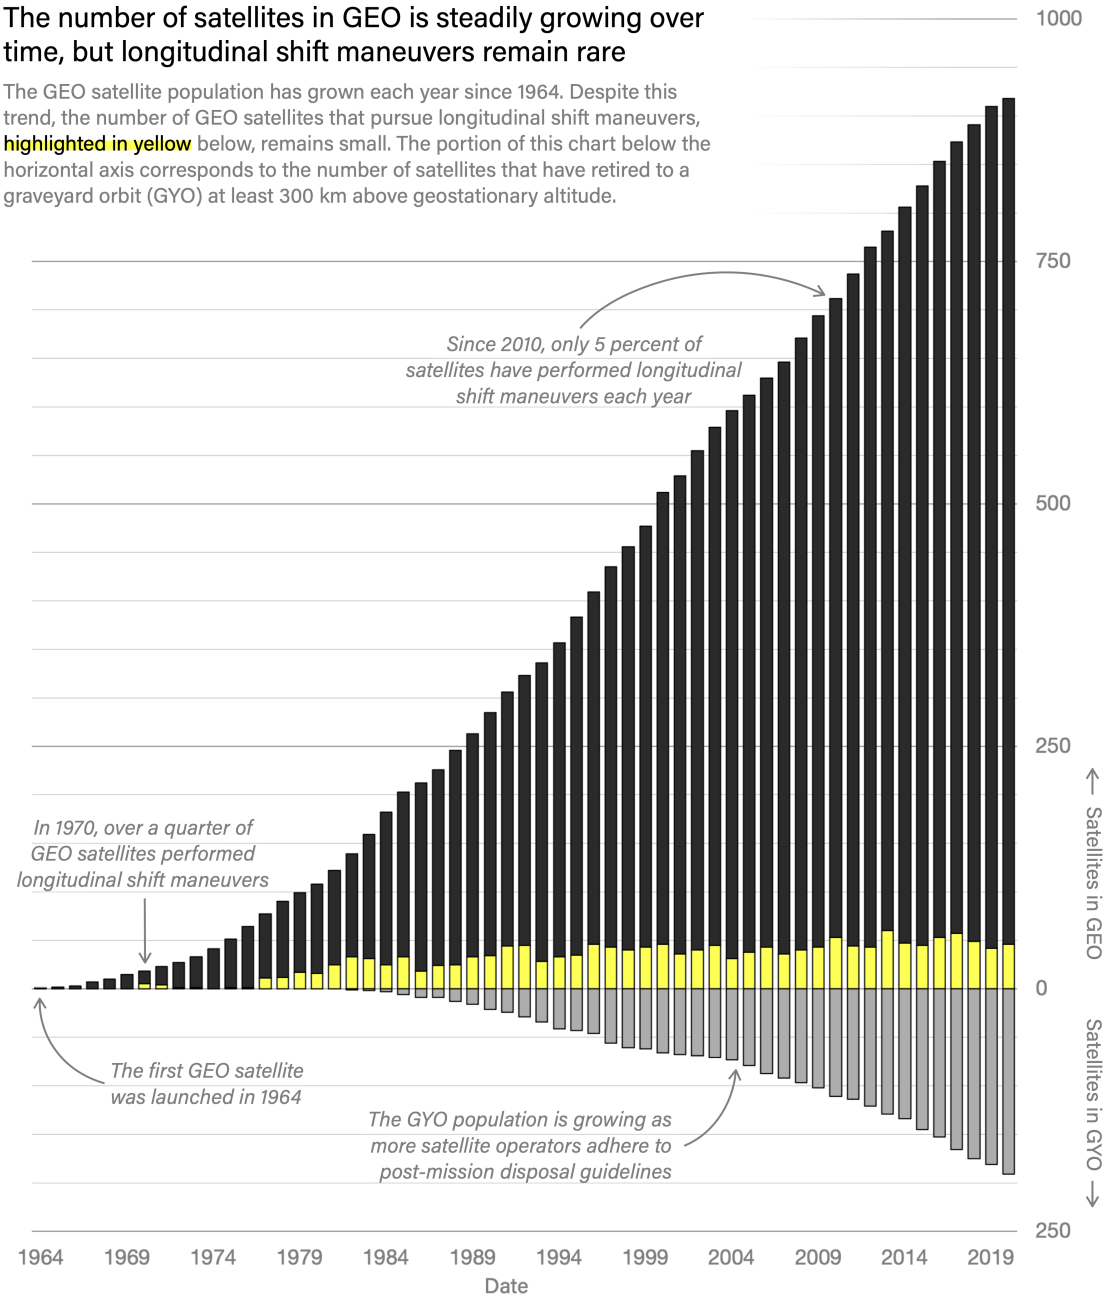

Fig. 1: Evolution of the GEO satellite population. The portion of this figure above the horizontal axis refers to satellites in GEO, as defined in Subsec. 2.1. The portion of the figure below the horizontal axis refers to satellites that were once in GEO, but have since spent the majority of a calendar year in a graveyard orbit (GYO), at least 300 km above geostationary altitude. Satellites that performed at least one longitudinal shift maneuver (either an IE-EE or IW-EW pair of maneuvers or a JE or JW maneuver, as defined in Subsec. 1.2) are highlighted in yellow.

to GEO has led to more satellites pursuing longitudinal shift maneuvers, satellites that perform longitudinal shift maneuvers remain greatly outnumbered by those that do not, as shown in Fig. 1. Over the past ten years, from 2011 to 2020, just 5 percent of satellites each year performed at least one longitudinal shift maneuver. Because longitudinal shift maneuvers are relatively rare, the maneuver detection algorithms described in Sec. 2 are specifically tailored to perform well on imbalanced datasets.

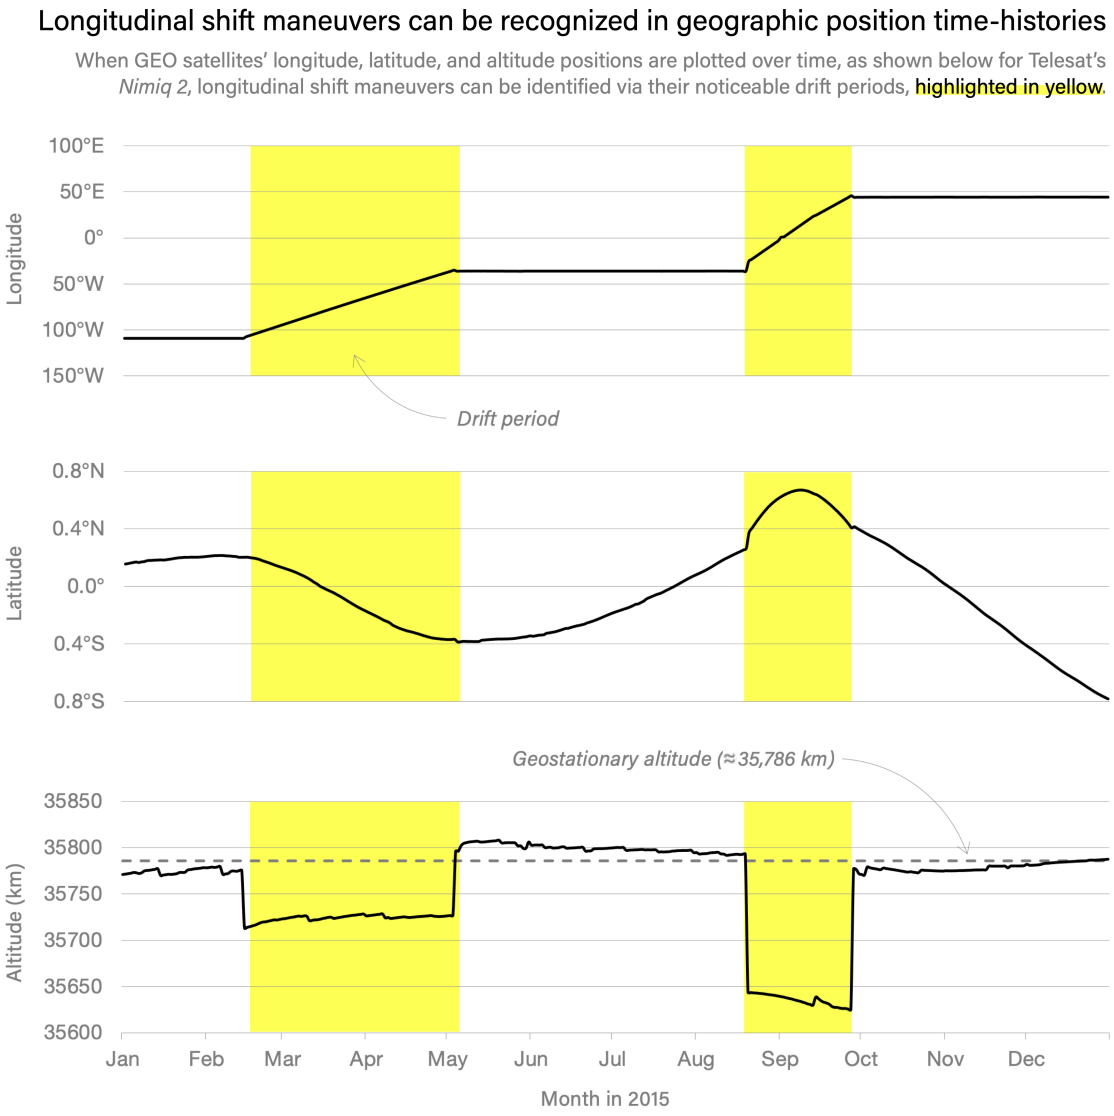

Fig. 2: Geographic position time-history for Telesat's *Nimiq 2* (Satellite ID: 27632) from January 1 to December 31, 2015. During this time, the satellite performed two longitudinal shift maneuvers, corresponding to two drift periods.

1.1 Identifying Maneuvers in Time-Series Geographic Position Data

Unlike satellite maneuvers in other orbital regimes, which are often identified through discontinuities observed over time in one or more of the classical orbital elements—eccentricity, semi-major axis, inclination, and other parameters—longitudinal shifts in GEO can be detected by human inspection of time-series geographic position data.

Fig. 2 describes the geographic position of a GEO satellite from January 1 to December 31, 2015. The first subplot, which depicts the satellite's longitudinal position over time, clearly shows that it began the year performing station-keeping maneuvers to maintain a sub-satellite longitude of approximately 109.1°W, an ideal position for offering coverage of the continental United States and Canada. Then, in mid-February, it initiated a longitudinal shift maneuver, which resulted in an eastward drift period that lasted until early May, highlighted in yellow in the figure. During this drift period, the satellite orbited the Earth at a significantly lower orbital altitude than during the prior station-keeping period, as shown in the third subplot of Fig. 2. Recall the equation of a satellite's velocity in a circular Earth orbit, v,

as a function of its altitude, z [5]:

$$v = \sqrt{\frac{\mu_\oplus}{R_\oplus + z}} \tag{1}$$

where $\mu_\oplus$ is the Earth's gravitational parameter ($\approx$398,600 km^3/s^2) and $R_\oplus$ is the Earth's radius ($\approx$6,378 km). When z decreases, v increases: a satellite's orbital velocity is faster during an eastward drift period than during station-keeping at one particular longitudinal position. Since satellites in GEO traditionally orbit the Earth in an eastward direction—known as *prograde* motion—an increase in orbital velocity results in an eastward drift in geographic longitude, or a positive slope in the first subplot of Fig. 2. At the end of the first drift period, the satellite performed a second maneuver to end the eastward drift, and resumed station-keeping at a longitudinal position of approximately 36.0°W, with a sub-satellite point near the eastern coast of Brazil. The satellite maintained this position via station-keeping until mid-August, when it performed another longitudinal shift maneuver and initiated a second eastward drift. The satellite then performed a maneuver to end the eastward drift in late September, at which point it resumed station-keeping at a longitudinal position of 44.5°E over east Africa and the Arabian peninsula.

Although the satellite's altitude over time offers an explanation for the changes observed in its longitude over time, its longitudinal position history alone is sufficient to identify the beginning and end of each longitudinal shift maneuver. Similarly, the latitudinal position history can be used to confirm longitudinal shift maneuvers, but is not critical—and often not sufficient—to discern such maneuvers when longitudinal position histories are also available.

1.2 Components of a Longitudinal Shift Maneuver

Both of the longitudinal shift maneuvers shown in Fig. 2 are made up of two components associated with initiating and ending an eastward drift, which will henceforth be abbreviated "IE" and "EE," respectively. Similarly, a westward longitudinal shift is often made up of two components: initiating and ending a westward drift ("IW" and "EW").

Although both of the drift periods featured in Fig. 2 were long in duration when compared to the data's daily sampling rate, that is not the case for all longitudinal shifts maneuvers. When satellites are observed to change their longitudinal position with shorter drift periods, the beginning and end of the drift may be difficult to discern upon inspection. In these cases, the satellites appear to "jump" eastward or westward over time. Fittingly, these types of longitudinal shift maneuvers can be abbreviated "JE" and "JW" for eastward and westward jumps, respectively.

Tab. 1 describes the six components of longitudinal shift maneuvers discussed along with their abbreviations.

1.3 GEO Satellites Patterns of Life

Satellite patterns of life (PoLs) are descriptions of on-orbit behavior made up of sequences of both natural and non-natural behavioral modes, including routine station-keeping, other on-orbit maneuvers, and uncontrolled motion [7]. After a GEO satellite is inserted into the geostationary belt, its PoL may include any of the six components of longitudinal shift maneuvers described in Tab. 1 separated by periods of station-keeping and natural drift. Although PoLs can take many forms, some trends can be observed within the GEO satellite population.

Table 1: The six components of longitudinal shift maneuvers. While some longitudinal shift maneuvers are made up of just one of the components listed below (JE and JW), others are made up of a pair (either IE and EE or IW and EW).

Maneuver Type	Abbreviation	Description
Initiate Eastward Drift	**IE**	The maneuver associated with beginning a longitudinal drift with a positive slope that exceeds a specified duration threshold
End Eastward Drift	**EE**	The maneuver associated with ending a longitudinal drift with a positive slope that exceeded a specified duration threshold
Jump Eastward	**JE**	The maneuver associated with beginning and ending a longitudinal drift with a positive slope within a specified duration threshold
Initiate Westward Drift	**IW**	The maneuver associated with beginning a longitudinal drift with a negative slope that exceeds a specified duration threshold
End Westward Drift	**EW**	The maneuver associated with ending a longitudinal drift with a negative slope that exceeded a specified duration threshold
Jump Westward	**JW**	The maneuver associated with beginning and ending a longitudinal drift with a negative slope within a specified duration threshold

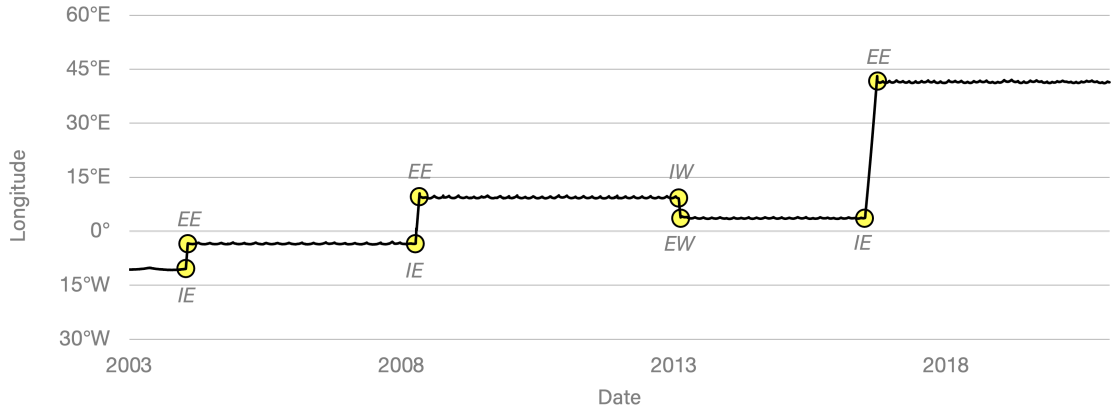

Fig. 3: Longitudinal position history for EUMETSAT's *Meteosat 8* (Satellite ID: 27509) from January 1, 2003, to December 31, 2020, annotated with the various maneuver types described in Tab. 1.

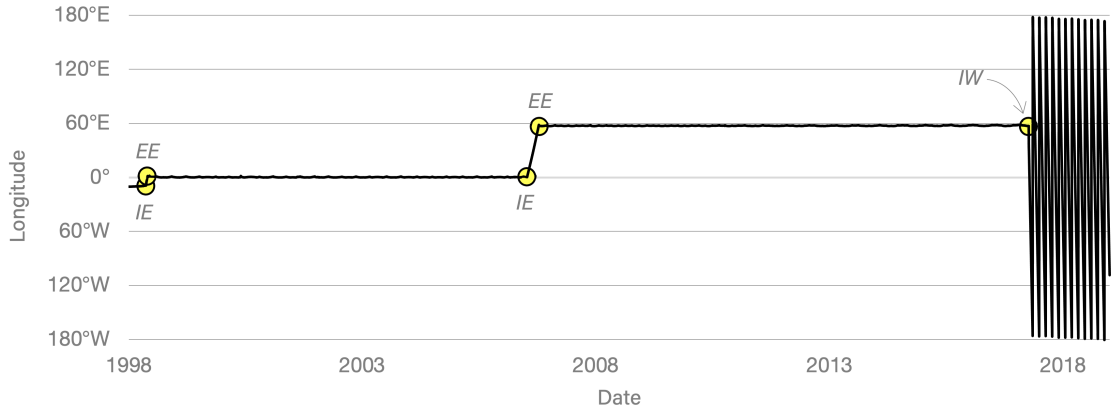

Fig. 4: Longitudinal position history for EUMETSAT's *Meteosat 7* (Satellite ID: 24932) from January 1, 1998, to December 31, 2018. Before it retired, *Meteosat 8* (shown in Fig. 3) performed a longitudinal shift maneuver to become closer to *Meteosat 7* and limit the coverage gap created by its retirement [21].

One of the most common PoLs for satellites in GEO is depicted in Fig. 3. Satellites that adhere to this PoL station-keep at various longitudinal positions, or *stations*, for years on end, and perform infrequent longitudinal shift maneuvers.

In an attempt to adhere to several U.S. and international guidelines, many satellite operators choose to remove their GEO satellites from the geostationary belt at the end of their operational lifetimes and place them in higher-altitude disposal orbits to reduce the risk of collision with other objects in GEO [16]. Satellites that raise both their apogee and perigee more than 300 km above geostationary orbit are said to have reached a *graveyard orbit* (GYO). The unique patterns associated with GEO satellite disposal are a key feature of another common satellite PoL in the geosynchronous region, depicted in Fig. 4.

Although most satellites in GEO remain operational for many years, others suffer from malfunctions that affect their ability to station-keep or perform longitudinal shift maneuvers. In those cases, satellite operators often choose to either abandon their satellites adrift in the geostationary belt or pursue a last-ditch maneuver to place them in a higher-altitude disposal orbit. Because malfunctions leave recognizable signatures in GEO satellites' geographic position time-histories, they represent yet another GEO satellite PoL, depicted in Fig. 5.

In addition to the PoLs highlighted thus far, others feature changes in geographic coordinate positions over time that are *not* due to any of the longitudinal shift maneuver components highlighted in Tab. 1, but rather changes in station-keeping patterns over time, including abandoning station-keeping in the east-west direction, north-south direction, or both. Fig. 6 shows the longitudinal position history for a satellite that abandoned its station and entered a *libration*

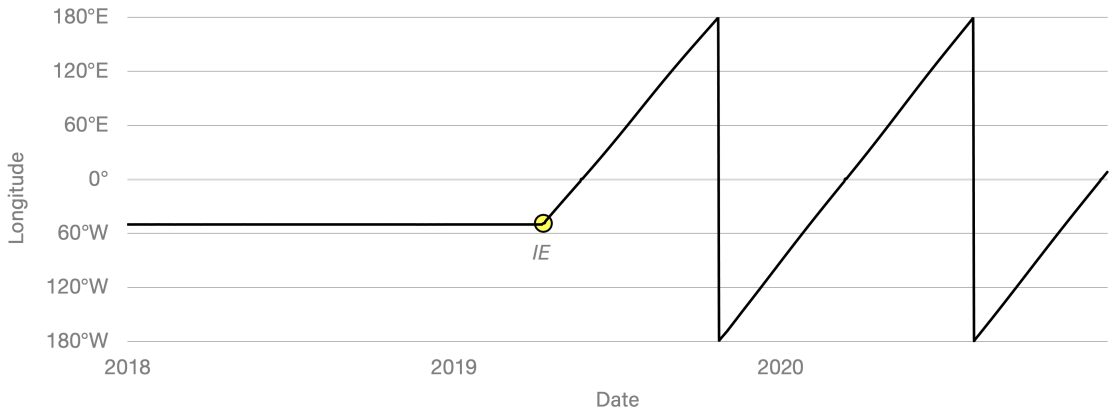

Fig. 5: Longitudinal position history for Intelsat's *Intelsat 29e* (Satellite ID: 41308) from January 1, 2018, to December 31, 2020. In 2019, the satellite developed a fuel leak followed soon thereafter by a loss of command and control [13].

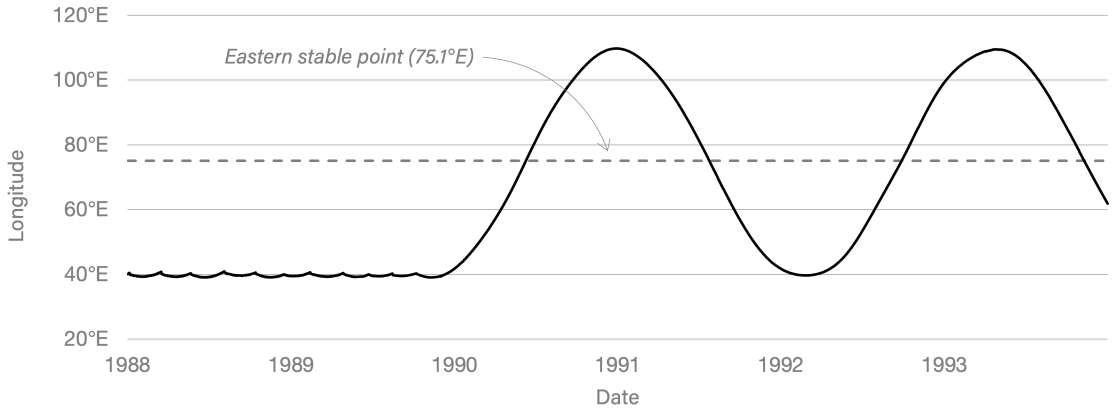

Fig. 6: Longitudinal position history for Russia's *Gorizont 3* (Satellite ID: 11648) from January 1, 1988, to December 31, 1993. In 1989, the satellite abandoned its station and entered a libration orbit around the eastern stable point.

orbit around the eastern stable point, due to the asymmetrical distribution of the Earth's mass [34, 71]. Libration orbits are not uncommon, especially for older satellites in GEO. In 2018, the European Space Agency (ESA) released a report tabulating the number of GEO objects, including debris, in various libration orbits [8]. The report identified 122 objects in libration orbit around the eastern stable point, 48 objects in libration orbit around the western stable point, and 19 objects in libration orbit around both stable points.

1.4 Previous Research on Satellite Maneuver Detection

Various methods for satellite maneuver detection have previously been developed that vary in their methodology, data input type and source, dependence on machine learning techniques, and effectiveness in various orbital regimes. This subsection briefly reviews several contributions to the literature, with a focus on the similarities and differences between them and the method described in the following section.

One popular strategy for satellite maneuver detection is centered on searching for statistical outliers in space objects' historical orbital element data without depending on machine learning techniques. Kelecy et al. [17] developed a method using smoothed orbital element time-series data generated from historical TLEs. In their paper, Kelecy et al. segmented the smoothed orbital element time-series data and fit polynomials to the data included in each segment. They then developed an algorithm that searched the entire smoothed orbital element dataset for relatively large differences between the polynomial fit value at the end of one segment and the beginning of the next. The algorithm used satellites' historical mean motion, semi-major axis, and orbital energy to detect in-plane maneuvers and inclination to detect out-of-plane maneuvers. Using tuned parameters, it could successfully detect maneuvers in

LEO satellite data at small Δv magnitudes on the order of several centimeters-per-second. Patera [25] developed a maneuver detection algorithm with the same premise as Kelecy's, but instead used a moving window as opposed to a segmentation strategy. Patera's method used orbital energy data derived from TLEs, and could be applied to satellites in any orbital regime.

A second popular strategy for satellite maneuver detection relies on propagating historical orbital parameters and comparing them to subsequent truth data. Relatively large differences between the predicted and actual parameter values are then flagged as maneuvers. This method was applied to maneuver detection in LEO by Lemmens and Krag [20] and Li et al. [22] using TLEs with promising results.

Newer studies apply both supervised and unsupervised machine learning techniques to satellite maneuver detection. Shabarekh et al. [33] presented a method for probabilistically characterizing GEO satellites' PoLs via unsupervised learning and predicting future maneuvers via supervised learning. The method was highly successful at learning and predicting station-keeping maneuvers in GEO, but was not applied specifically to longitudinal shift maneuvers. The study also did not depend on TLE data, but rather synthetic, relatively sparse orbital element data produced by the U.S. Air Force Research Lab. More recently, Bai et al. [2] applied a suite of unsupervised machine learning techniques—including K-means and hierarchical cluster analysis—to historical TLE data and classified orbital maneuvers in LEO by magnitude.

The method described in the following chapter was first formulated by Roberts and Linares [31]. This framework primarily differed from those described earlier in two ways: it relied on historical geographic position data as opposed to classical orbital elements and used a CNN algorithm design. While the method described in the following chapters also depends on historical geographic position data and CNNs, it uses a more refined data labeling strategy, and applies the developed algorithms' results to GEO satellite anomaly detection. In the previous study, longitudinal shift maneuvers were labeled in their entirety, without regard to the individual maneuver components described in Tab. 1 of which they are composed. For example, an eastward shift maneuver like the ones described in Fig. 2 would be labeled as one long longitudinal shift maneuver as opposed to two smaller ones. In the refined labeling strategy, the same maneuver would be labeled with more detail: the beginning of the eastward drift would be labeled as an IE maneuver, the drift itself would remain unlabeled, and the end of the drift would be labeled as an EE maneuver. The methodology described in the next section uses the six longitudinal shift maneuver components described in Tab. 1 to train and test six different satellite maneuver detection algorithms, which together work as a suite for maneuver classification. This method allows for more consistent maneuver labeling and leads to more precise results, as discussed in Section 3. Additionally, the suite of maneuver detection algorithms is trained on more historical data, improving its recall when compared to the algorithm evaluated in [31].

2. METHODOLOGY

This study is focused on the on-orbit behavior of active satellites in the geosynchronous region. The primary data source for this work is a space object catalog provided by the United States Space Force's (USSF) 18th Space Control Squadron (18 SPCS)—the space control unit responsible for managing the U.S. Space Command's (SPACECOM) SSA program—and made available online to the public at Space-Track.org [35]. Historical orbital parameters are described in the catalog using TLEs, which encode space objects' orbits at particular times, known as *epochs*. A collection of TLEs, ordered by epoch, is known as a *TLE time-history* [17].

2.1 Defining the GEO Region and Identifying GEO Satellites

Although the SPACECOM catalog includes various types of space objects—such as payloads, rocket bodies, and debris—this analysis focuses only on the behavior of active satellites, and thus only references data for payload objects.[1] Since this study is focused on the geosynchronous region, only payload objects that were determined to be in GEO at least once before the end of the study period on December 31, 2020, are referenced. In this case, a geosynchronous orbit is defined as any orbit that is fully enclosed in the geosynchronous region as described by the Inter-Agency Space Debris Coordination Committee (IADC) [15].[2]

[1] Objects that were uncategorized in the catalog on the date it was accessed (February 1, 2021), were also included in this study, since those objects may be re-labeled as payloads once they are categorized by 18 SPCS.

[2] By this definition, GEO orbits have both an apogee and perigee altitude within 200 km of geostationary altitude and an inclination between -15 and 15 degrees. Although some objects may pass through the geosynchronous region during their orbital periods, this definition for GEO only includes objects that are in this region for their entire orbital period.

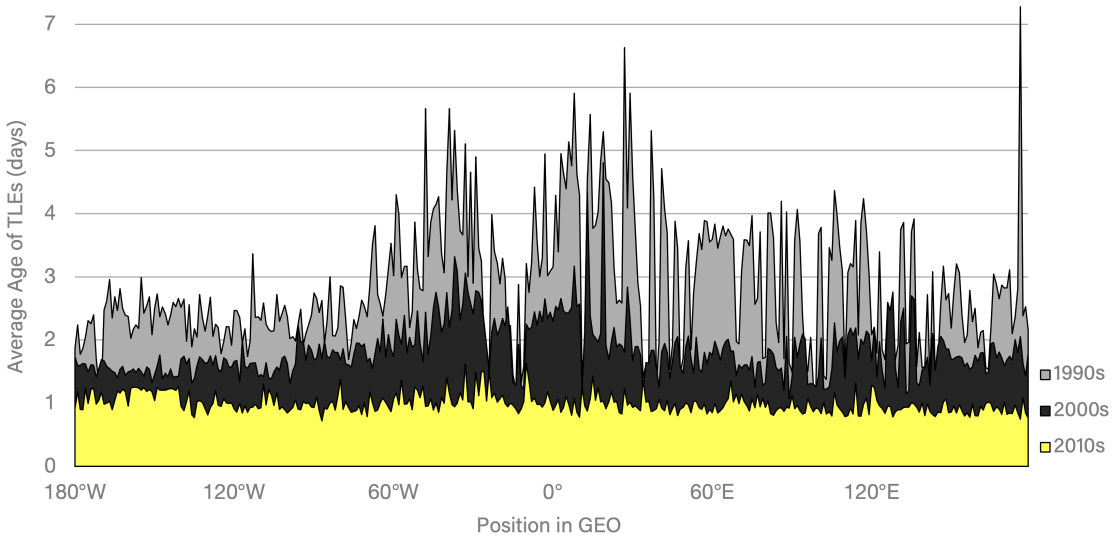

Fig. 7: Average age of TLEs over time for GEO satellites. In general, the average age of TLEs has decreased over time. In the 2010s, TLEs satellites at all longitudinal positions are published approximately once per day on average.

To determine whether an individual TLE places its corresponding object in GEO using the IADC's definition, a portion of the TLE's encoded orbital elements must be converted into apogee and perigee altitudes (h_a and h_p, measured in meters above the Earth's surface) using Eq. 1 and Eq. 2 from [31]. After performing this calculation on the TLEs for all payload objects available in the SPACECOM catalog, it can be determined that there have been 1,109 satellites in GEO for at least one day during the study period.[3]

2.2 Converting TLE Time-histories to Geographic Coordinates

The geographic coordinate data presented in Fig. 2, also known as a *geographic time-history*, is evenly spaced in time. TLE time-histories, however—which are used to derive geographic time-histories—are not evenly spaced in time. The epochs encoded in each TLE time-history are more densely spaced for some satellites and more sparsely spaced for others. To convert unevenly spaced TLE time-histories to evenly spaced geographic time-histories, an appropriate time step must first be selected at which to sample the TLE time-histories. From 2010 to 2020, the average separation between epochs for GEO satellites, also known as the *average age of TLEs* was approximately 1 day for all longitudinal positions, as shown in Fig. 7. Thus, a one-day time step is an appropriate choice for this study.

At each time step in a satellite's geographic time-history, one TLE must be selected from the satellite's TLE time-history with which to calculate the satellite's geographic position. The appropriate TLE can be found by searching the satellite's TLE time-history and identifying the TLE with an epoch closest to a given time step in the geographic time-history, known as the *nearest TLE*. If the nearest TLE's epoch is earlier than a given time step in a satellite's geographic time-history, the TLE must be propagated forward to the date and time associated with the given time step. Similarly, if the nearest TLE's epoch is later than the given time step, the TLE must be propagated backwards in time. Since TLEs lose accuracy with time, TLEs should not be propagated forward or backwards by more than two weeks.[4] This study uses the SGP4/SDP4 model [37] included in the PyEphem Python library for propagation [27]. The propagated TLE, known as the *time-adjusted nearest TLE*, can then be saved as a PyEphem Object.

At each time step, the orbital parameters encoded in the time-adjusted nearest TLE can be converted into geographic coordinates—latitude, longitude, and altitude—algorithmically using a coordinate system transformation [36, 172]. This study uses the PyEphem Object's built-in "sublat," "sublong," and "elevation" properties to compute the geographic positions directly from the time-adjusted nearest TLE.

Recall from the previous section that not all 1,109 GEO satellites included in this study necessarily spend the entire study period within the geosynchronous region. Although all of the GEO satellites spend at least one day in the GEO

[3] This figure can be seen as the sum of the satellites above and below the horizontal axis in the 2020 portion of Fig. 1.

[4] If the nearest TLE's epoch is more than two weeks away from a given time step, no geographic position should be calculated for that time step.

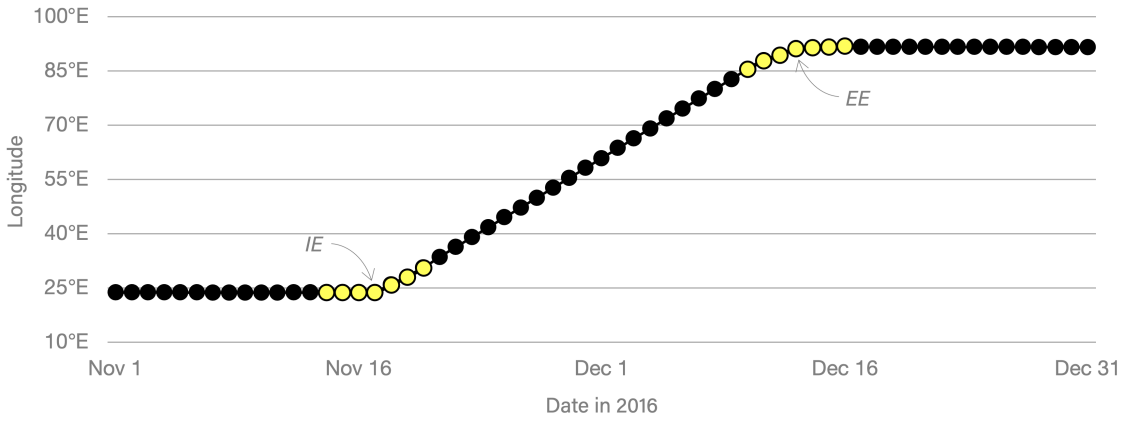

Fig. 8: An example of the labeling process for two types of longitudinal shift maneuvers. This figure shows daily longitudinal positions for Telesat's *Nimiq 2* from November 1 to December 31, 2016. Two components of a longitudinal shift maneuver are labeled in this plot with a maneuver buffer period of 3 days.

region according to Subsec. 2.1's definition of a GEO satellite, they each spend a portion of their operational lifetime below GEO (in the days after they are launched, but before they reach GEO) and many of them leave GEO once reaching it (perhaps due to a longitudinal shit maneuver, an on-orbit malfunction, retirement to GYO, or abandoning station-keeping). To record whether each object falls within the GEO region at each time step in the study period, a new binary field called "GEO Status" can be included in the geographic position time-history as a fourth feature after geographic latitude, longitude, and altitude. Since this study is focused on satellite behavior in GEO (as opposed to anything that happens before the satellite reaches GEO), the geographic time-histories should be truncated such that the first time step in the time-history corresponds to the first day in which the satellite falls within the GEO region.

2.3 Labeling Longitudinal Shift Maneuvers

The longitude portion of geographic time-histories can be plotted and the six types of longitudinal shift maneuvers featured in Tab. 1 can be labeled by inspection. In addition to labeling any time steps on which maneuvers occur, a number of time steps before and after maneuvers, known as the *maneuver buffer period* (b), should also be labeled. Including a maneuver buffer period of any size will assist the CNN described later in this section in recognizing the various maneuver types, as maneuver patterns are most recognizable when observed alongside geographic position data from before and after the day the maneuver occurs. Fig. 8 shows an example of manually labeled longitudinal position data with a buffer period of 3 days. If a drift period is too short to accommodate the buffer periods associated with a IE-EE or IW-EW pair of maneuvers—that is, its length is less than or equal to $2b+1$ days—then that drift period should be labeled JE or JW, respectively, with a buffer period labeled both before and after.

Under the described labeling strategy, an individual day should only be labeled with one of the six longitudinal shift maneuver types described in Tab. 1, but consecutive days may be labeled with two different types. Additionally, any individual maneuver type should appear on no fewer than $2b+1$ consecutive days, as shown in Fig. 8.

2.4 Segmenting the Geographic Time-histories

In order to create a dataset on which to train the maneuver detection algorithms to identify the various maneuver types, each geographic time-history—regardless of its length—should be divided into segments and each segment labeled with the types of maneuvers it contains, if any.

To segment the geographic time-histories, a frame width (f) and step size (s) should be selected, such that the segments are each $f+1$ days long and their start dates are separated by s days. Notice that when $s \leqslant f$, individual days appear in more than one segment. When $f = 28$ and $s = 14$, for example, each date appears in two different segments, which will later offer the detection algorithms two opportunities to detect maneuvers on each day.

After the segments are created, they must be labeled based on whether any of the maneuver types from Tab. 1 appear during that period. Since maneuver types are most recognizable when they are preceded and followed by a buffer period, only segments that contain both maneuvers and their corresponding buffer periods in their entirety should be labeled as such. That is, a segment should be labeled with a maneuver type if it contains no fewer than $2b+1$

individual days with a particular maneuver label. To develop a training dataset for the maneuver detection algorithms described in the next section, the frame width, step size, and maneuver buffer period should be selected such that $f \geqslant s$ and $s > 2b + 1$. These considerations increase the number of opportunities the maneuver detection algorithms have to detect maneuvers and ensure that at least some segments contain labeled maneuvers.

Note that under the described labeling strategy, an individual segment may be labeled with more than one of the maneuver types described in Tab. 1.

2.5 Preprocessing the Geographic Time-history Segments

Prior to training the maneuver detection algorithms, the geographic time-history segments should be preprocessed to account for two peculiarities of the longitudinal measurement system.

The geographic time-history segments, like the geographic time-histories themselves, measure longitude on a scale of -180 to 180 degrees (corresponding to 180°W and 180°E, respectively). Despite this convention, there is no physical significance of the 180th meridian in this measurement system. When a satellite traverses the 180th meridian, however—meaning it crosses from 179.9°E to 179.9°W, or vice versa, due to a longitudinal shift maneuver or as part of the satellite's natural drift—it may appear to an untrained observer like a discontinuity on the longitudinal position plot. To account for this issue, geographic time-history segments that include 180th meridian crossings should have their longitudinal scales adjusted. For example, if a satellite begins a segment with a longitudinal position of 175°E, then initiates and ends an eastward drift such that it ends the segment with a longitudinal position of 175°W, that segment's longitudinal scale should be adjusted: instead of ranging from -180 to 180 degrees, it should range from -180 to 540 degrees, and all western hemisphere measurements should be increased by 360 degrees. Such a change would make the example segment appear to begin at a longitude reading of 175 degrees and end at 185 degrees, with no observable discontinuity at the 180th meridian.

In addition, the geographic time-history segments should be normalized such that their minimum and maximum longitudinal measurements are mapped to 0 and 1 respectively. This process allows the maneuver detection algorithms to be trained to detect maneuvers regardless of their starting and ending positions on the longitudinal scale and the drift rate associated with any shift maneuvers.

2.6 Resampling the Training Dataset

In order to measure the success of the maneuver detection algorithms designed in the following section, a portion of the geographic time-history segments must be reserved and not used for training. For this study, all segments that contain a date in the year 2020 are reserved for this purpose, and are known as the *test dataset*. The segments that contain dates between 2010 and 2019 form the *train dataset*, will be used to train the algorithms.

Since the six maneuvers described in Tab. 1 are relatively rare, the vast majority of the geographic time-history segments used for training contain no maneuvers of any kind. To obtain a more balanced training dataset—where the number of segments that contain maneuvers (known as the *positive class*) is closer to the number that contain no maneuvers (known as the *negative class*)—resampling strategies should be used on the train dataset to balance the two classes [3]. When used in combination, over-sampling the positive class (using methods such as SMOTE [4] or ADASYN [11]) and under-sampling the negative class (by simply removing a portion of the segments that contain no maneuvers) have been measured to perform particularly well [12].

2.7 Algorithm Design

Six one-dimensional convolutional neural networks (CNNs) can be trained to identify each of the six types of longitudinal shift maneuvers described in Tab. 1 using the Keras Python library [9]. This subsection describes the parameters that were selected to create and segment all GEO satellites' geographic time-histories and form the dataset needed to train the CNN-based algorithms. It also includes a description of the parameters selected to develop preliminary versions of the algorithms for each of the maneuver types that are evaluated in the following section. Note that the following selections represent an initial design for a suite of maneuver detection algorithms, not the result of a hyperparameter optimization process. Future studies would likely benefit from tuning the hyperparameters, which may lead to improvements in precision and recall when compared to the results presented in the following section.

The algorithms evaluated in the following section depend on geographic time-histories labeled using the labeling strategy described in the Subsec. 2.3. To create the geographic time-history segments, maneuvers were labeled with a 3-day maneuver buffer period ($b = 3$). Such a selection ensures that each maneuver is labeled within a cluster of at least 7 days, as shown in Fig. 8. After the geographic time-histories were labeled, they were divided into segments of

length 28, with a 1-day separation between the start dates of the segments ($f = 28$; $s = 1$). These selections ensure that each day within a GEO satellite's geographic time-history appears in 29 different segments, offering the algorithms many opportunities to identify maneuvers on each day. After the geographic time-history segments were created, the positive class within the training dataset was over-sampled using the ADASYN method with a 0.75 sampling strategy and the negative class was randomly undersampled with a 0.80 sampling strategy.

This re-sampled training data was then used to fit a CNN model with the following layers: a one-dimensional CNN layer (created using Keras' "Conv1D" class) with 32 filters, a kernel size of 3, and the rectified linear unit (ReLU) activation function; a second one-dimensional CNN layer with the same parameters as the first layer; a dropout layer (created using Keras' "Dropout" class) with a frequency of 0.5; a one-dimensional max pooling layer (created using Keras' "MaxPooling1D" class) with a pooling size of 2; a dense layer (created using Keras' "Dense" class) with a size of 100 and the ReLU activation function; and a second dense layer with the "Softmax" activation function, sized to the number of geographic time-history segments included in the training dataset.

In the following section, each of the six maneuver detection algorithms designed in this section are evaluated after being trained on ten year's worth of data (corresponding to dates between January 1, 2010, and December 31, 2019) and tested against geographic time-history segments from January 1 to December 31, 2020.

3. RESULTS

The algorithms described in the previous section can be used to detect each of the six satellite maneuver types using any GEO satellite's geographic position data that is segmented using the same segmentation strategy as the training dataset. This section evaluates the geographic time-histories segments for all satellites in GEO in 2020.

With a frame width of 28 days, a step size of 1 day, and hundreds of operational GEO satellites in orbit, there are more than 400,000 segments in this test period, of which less than 0.4 percent contain maneuvers.

3.1 Performance Evaluation Metrics

The following subsection describe the performance of each algorithm via the distribution of the four components of a binary confusion matrix [18, 254]: "True positive" (outcomes in which the algorithms detect a maneuver that did indeed occur in a particular segment), "False positive" (outcomes in which the algorithms detect a maneuver that actually did not occur in a particular segment), "False negative" (outcomes in which the algorithms fail to detect a maneuver that really did occur in a particular segment), and "True negative" (outcomes in which the algorithms do not detect a maneuver in a particular segment in which a maneuver did not occur), abbreviated TP, FP, FN, and TN, respectively.

The four components of the confusion matrix are useful for calculating critical metrics that can be used to evaluate the algorithms' performance, including *accuracy*, *recall*, *precision*, and the F_1 score.[5]

Table 2: Performance metrics for the suite of satellite maneuver detection algorithms on a segment-by-segment basis after they were trained on ten year's worth of GEO satellite data (geographic time-history segments containing dates between January 1, 2010, and December 31, 2019) and evaluated on the GEO satellite population in 2020.

	TP	FP	FN	TN	A	R	P	F_1
IE	400	983	215	398,952	99.7%	65.0%	28.9%	40.0%
EE	454	2,057	43	397,996	99.5%	91.3%	18.1%	30.2%
JE	115	592	63	399,780	99.8%	64.6%	16.3%	26.0%
IW	1,277	1,230	148	397,895	99.7%	89.6%	50.9%	65.0%
EW	691	1,208	63	398,588	99.7%	91.6%	36.4%	52.1%
JW	102	772	2	399,674	99.8%	98.1%	11.7%	20.9%

[5] Accuracy (A), precision (P), and recall (R) can be calculated as folows: $A = \frac{TP+TN}{TP+FN+TN+FP}$; $R = \frac{TP}{TP+FN}$; and $P = \frac{TP}{TP+FP}$. The F_1 score is the harmonic mean of the recall and precision: $F_1 = \frac{2\times TP}{2\times TP+FP+FN}$.

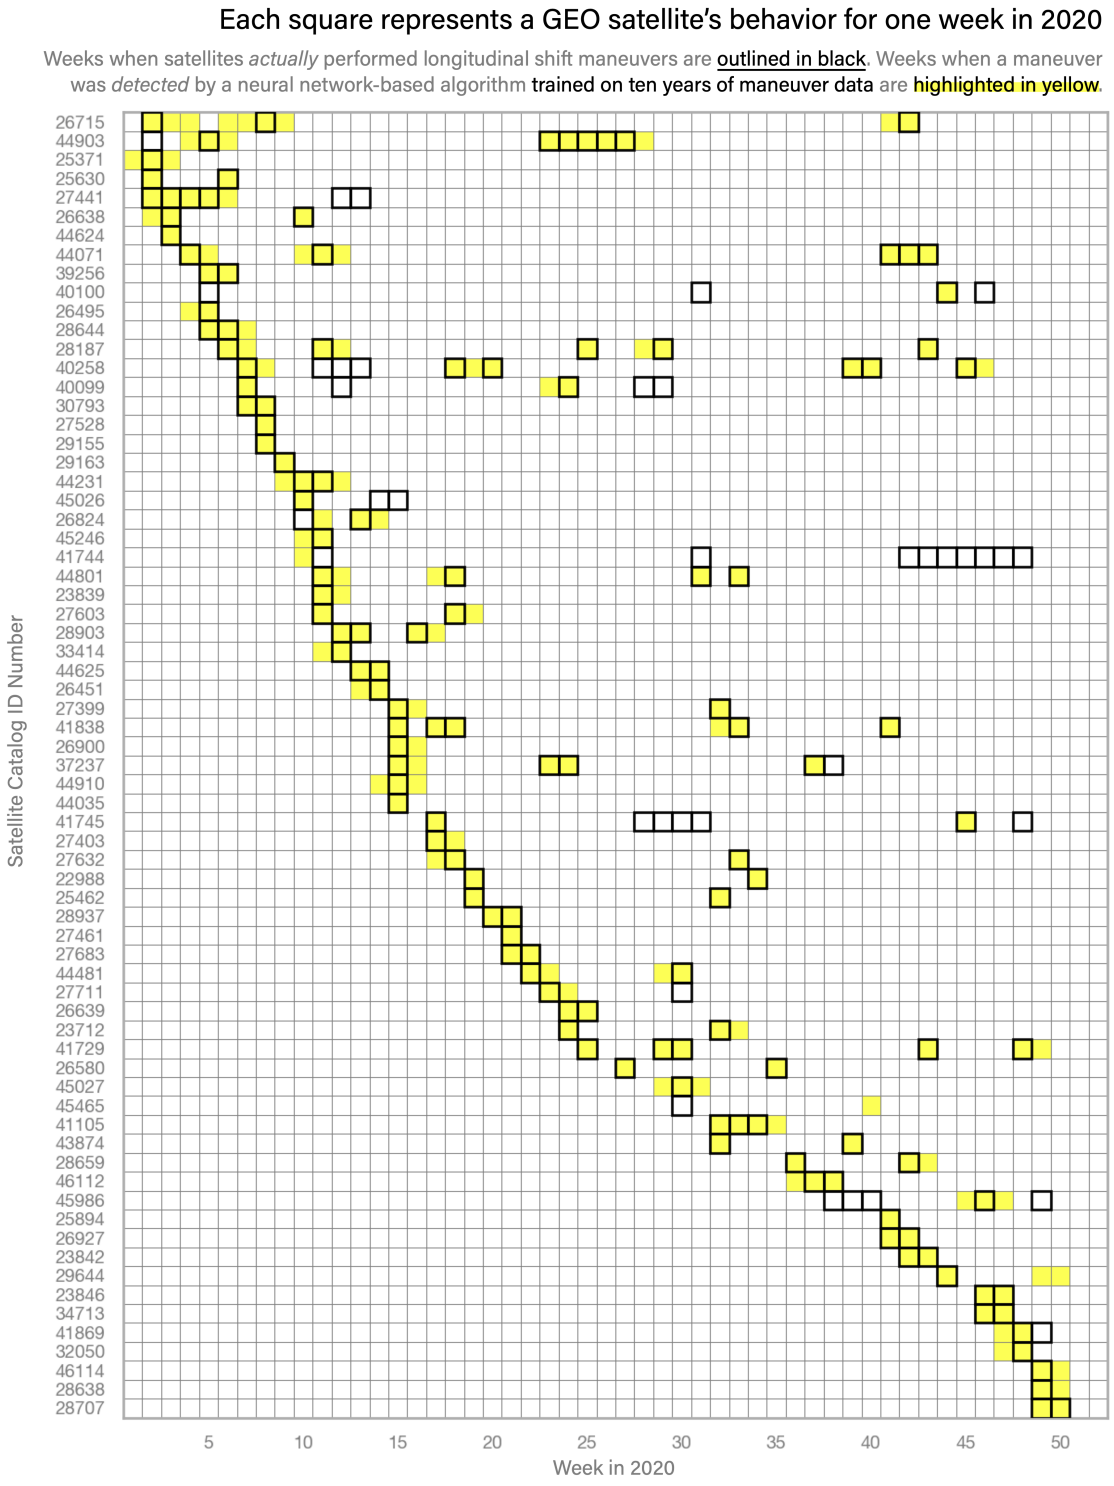

Fig. 9: Detection results for the suite of satellite maneuver detection algorithms on a maneuver-by-maneuver basis after training on ten years' worth of GEO satellite data. The rows of the figure correspond to the 71 GEO satellites that performed a longitudinal shift maneuver of any kind in 2020 ordered by the week of their first maneuver.

3.2 Algorithm Performance

Tab. 2 describes the four components of the confusion matrix for each of the six satellite maneuver detection algorithms after they were trained on ten years' worth of GEO satellites' geographic time-history segments. The last four columns of the table present the results of the performance metrics calculations described in the previous subsection.

All six maneuver types exhibited recall rates of approximately 65 percent or more. For four of the six types, those associated with ending a longitudinal drift, the recall was much higher: approximately 90 percent of the positive class or more was identified. All six of the algorithms, however, exhibited poor precision; false positives often outnumbered true positives. Since the F_1 scores weigh both recall and precision, they too were poor: only two of the six algorithms have an F_1 score over 50 percent.

Although Tab. 2 provides a detailed view of the algorithm suite's performance on a segment-by-segment basis, satellite operators and SSA analysts may be more interested in its performance on a maneuver-by-maneuver basis. Recall that each maneuver that actually occurred in 2020 is captured in more than one segment. A maneuver that is labeled over 7 days, for example, like the two shown in Fig. 8, appears fully in 23 different segments in this segmentation strategy. Instead of reporting the results of the algorithm on each of those 23 segments, consider their performance as a group.

Fig. 9 uses this concept to plot the suite of maneuver detection algorithm's performance on a maneuver-by-maneuver basis. Each row in the figure corresponds to a satellite that actually performed at least one type of longitudinal shift maneuver in 2020 and is labeled with the satellite's SPACECOM Satellite ID number. Each square represents a satellite's behavior for one week in 2020. Those outlined in black contain a maneuver—that is, four of the seven days in that week were *labeled* with one of the six different maneuver types. Squares in which any of the six algorithms identified a maneuver—that is, at least four days in that week were *predicted* to contain maneuvers (meaning those days' 23 segments were each flagged as containing a maneuver)—are highlighted in yellow.

The confusion matrix under this maneuver-by-maneuver formulation is more favorable than that of the segment-by-segment formulation for each maneuver type presented earlier in this section. With 1,104 satellites in the test dataset in 2020 and 52 weeks per year there are 57,408 total outcomes in the confusion matrix, with 173 in the positive class and 57,235 in the negative class and the following distribution: 122 true positives, 47 false positives, 51 false negatives, and 57,188 true negatives.

This confusion matrix corresponds to an accuracy of 99.8 percent, recall of 70.5 percent, precision of 72.2 percent, and an F_1 score of 71.3 percent.

4. DISCUSSION

When satellites reach the end of their operational lifetime, they are often removed from GEO and placed in GYO, as described in Subsec. 1.3 and shown in Fig. 1. Thus, the act of retiring a satellite may be considered nominal behavior for relatively old satellites, but anomalous behavior for younger ones. Fig. 10 shows a portion of a longitudinal position history for a satellite that was launched in May 2006 and performed no longitudinal shift maneuvers until its retirement in late February 2020. With almost fourteen years of station-keeping in GEO, the satellite's detected retirement could be considered nominal behavior. Two suites of algorithms—the one trained on ten years' worth of GEO satellite maneuver data described in Sec. 2 and a modified version of the same suite trained on just five years' worth of data from 2015 to 2019—were successful in detecting the IW maneuver associated with the retirement.

The pair of algorithm suites did not perform equally well, however, in identifying an anomalous retirement that appeared in the test dataset. Only the suite of algorithms trained on ten years' worth of GEO satellite maneuver data successfully detected the IW maneuver associated with the satellite's retirement featured in Fig. 11. Although this satellite retired nearly 15 years after its launch in April 2005—meaning it had a longer operational lifetime than other GEO satellites, including the one featured in Fig. 10—it did so in response to an on-orbit malfunction, not simply because it reached the end of its operational lifetime [14]. According to the operator, the satellite's onboard batteries suffered "significant and irreversible thermal damage" as part of an anomaly detected in December 2019 that led to "a significant risk that [the] battery cells could burst" [32]. After receiving permission from the U.S. Federal Communications Commission to skip the standard guidelines surrounding expending all remaining fuel at the end of a satellite's operational lifetime, the operator rapidly decommissioned the satellite on short notice [24]. The IW maneuver associated with the satellite's non-standard retirement was only detected by the suite of algorithms trained on ten years' worth of historical GEO maneuver data. The suite trained on only five years of data did not detect it.

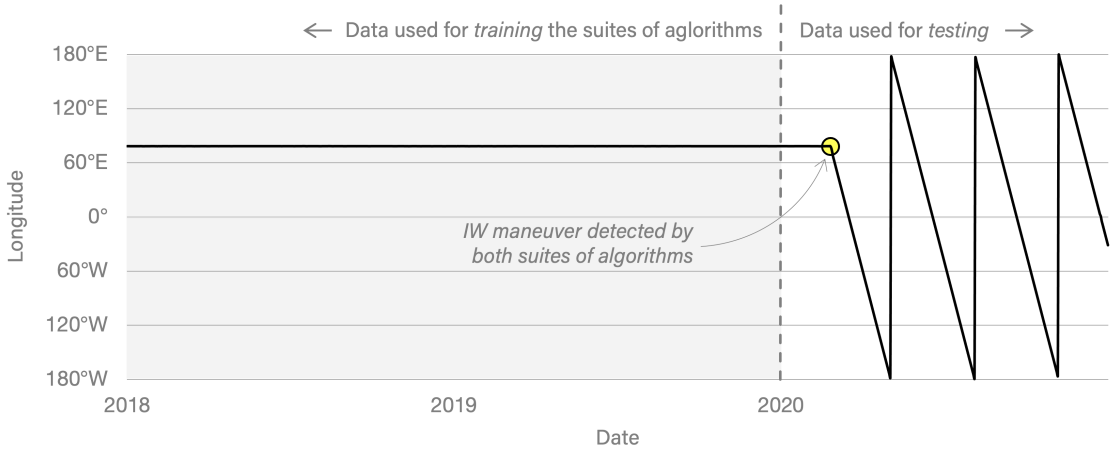

Fig. 10: Maneuver detection results for Thaicom's *Thaicom 5* (Satellite ID: 29163), including daily longitudinal positions from January 1, 2018, to December 31, 2020. This satellite was launched on May 27, 2006, and occupied approximately the same longitudinal position, 78.4°E, from July 2006 until its retirement in February 2020. Two suites of satellite maneuver detection algorithms—the one trained on ten years' worth of GEO satellite maneuver data and the one trained on ten—correctly detected the IW maneuver.

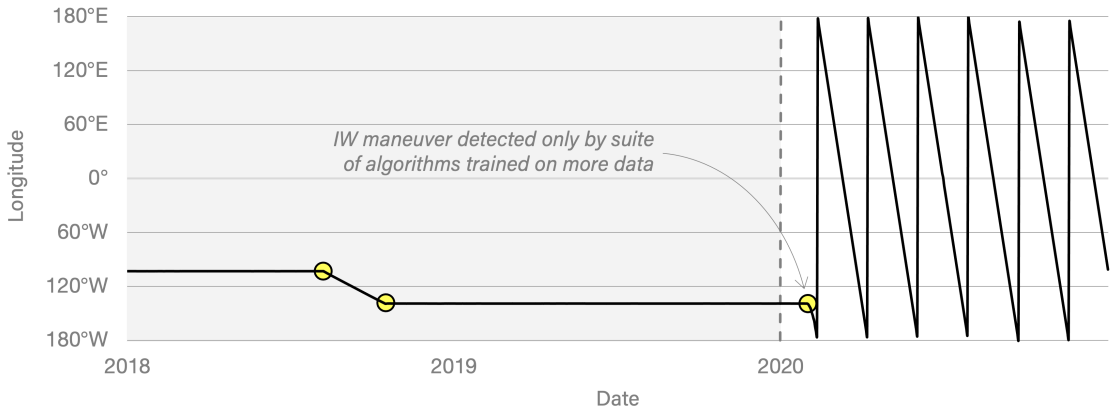

Fig. 11: Maneuver detection results for DirecTV's *Spaceway 1* (Satellite ID: 28644), including daily longitudinal positions from January 1, 2018, to December 31, 2020. This satellite was launched on April 26, 2005, and occupied one longitudinal position, 102.8°W, from August 2005 to August 2018, and a second position, 138.9°W, from October 2018 until its retirement in late January and early February 2020.

4.1 Identifying Nominal and Anomalous Longitudinal Shifts

When satellites perform a longitudinal shift maneuver, their previous PoL—the description of their behavior prior to the maneuver taking place—can help determine the nominal or anomalous nature of the maneuver.

The satellite featured in Fig. 12 has performed more longitudinal shift maneuvers than almost any other GEO satellite in the SPACECOM space object catalog. In 2020, it performed one westward shift maneuver. Although the maneuver featured a longer drift period than that of any previous longitudinal shift maneuver, the satellite's drift rate was approximately equal to those previous. Since the satellite's 2020 behavior largely adheres to its previous PoL, the newest maneuver can be considered nominal. Both the five- and ten-year satellite maneuver detection algorithm suites detected the IW and EW maneuvers performed in 2020.

The satellite featured in Fig. 13 performed no longitudinal shift maneuvers after its launch in April 2009 until November 2020. Because the satellite's 2020 behavior differs from its previous PoL, its westward longitudinal shift maneuver

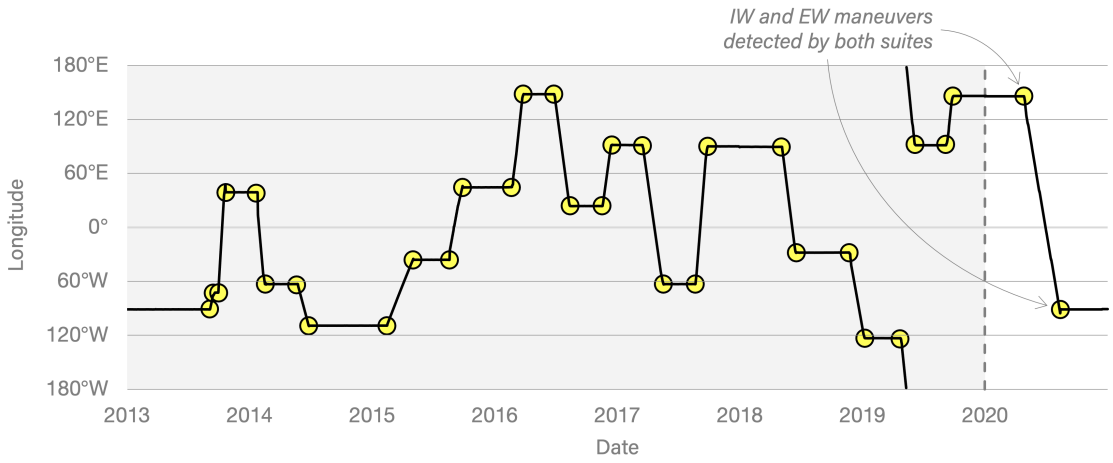

Fig. 12: Maneuver detection results for Telesat's *Nimiq 2*, including daily longitudinal positions from January 1, 2013, to December 31, 2020. The satellite was launched on December 29, 2002, and has occupied at least 16 longitudinal positions in GEO since then, making it one of the most maneuvered satellites in the SPACECOM catalog. In 2020, the satellite performed a longitudinal maneuver with a longer drift period than ever before, but a familiar drift rate.

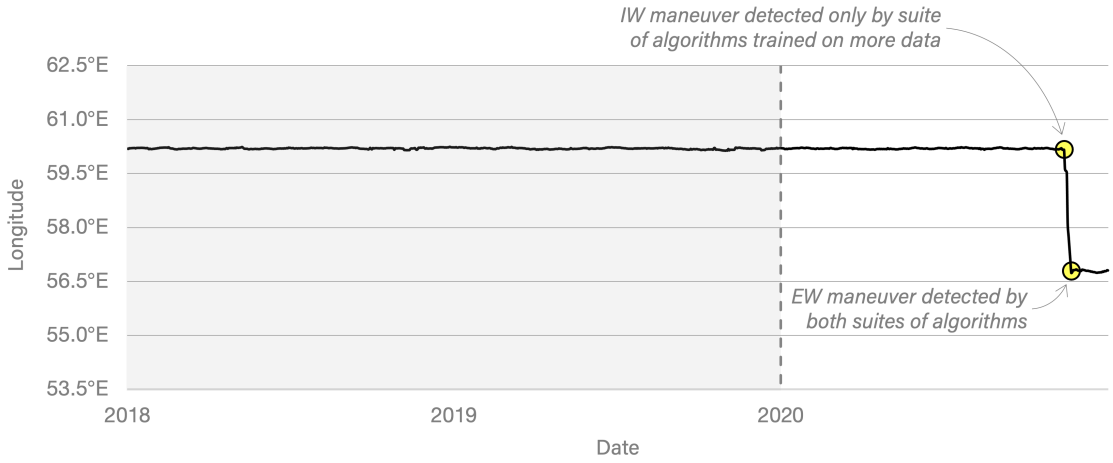

Fig. 13: Maneuver detection results for the U.S. Air Force's *WGS F2* (Satellite ID: 34713), including daily longitudinal positions from January 1, 2018, to December 31, 2020. The satellite was launched on April 4, 2009, and occupied approximately the same longitudinal position, 60.2°E, from the beginning of its geographic position time-history until it performed a westward longitudinal shift maneuver in November 2020.

could be considered anomalous. Unlike most maneuvers, however, the satellite's initial IW maneuver was difficult to detect by the suites of machine learning algorithms, possibly due to its relatively small longitudinal displacement: less than 4 degrees over the course of nine days. Although both the suite of algorithms trained on five years of historical data and ten years detected the EW maneuver associated with the end of the satellite's 2020 westward shift maneuver, only the suite trained on more data caught the IW maneuver.

4.2 Well-known Cases of Anomalous Behavior in GEO

Fig. 14 and Fig. 15 feature the longitudinal position histories for two satellites with exceptionally unique PoLs: Russia's *Luch (Olymp)* (Satellite ID: 40258) and China's *SJ-17* (Satellite ID: 41838).

Luch (Olymp) is a Russian GEO satellite that was launched in September 2014. Since then, it has occupied over 20 longitudinal positions in the geostationary belt—more than any other GEO satellite in the SPACECOM space object catalog—and approached several uncooperative satellites in the process, attracting attention from the international

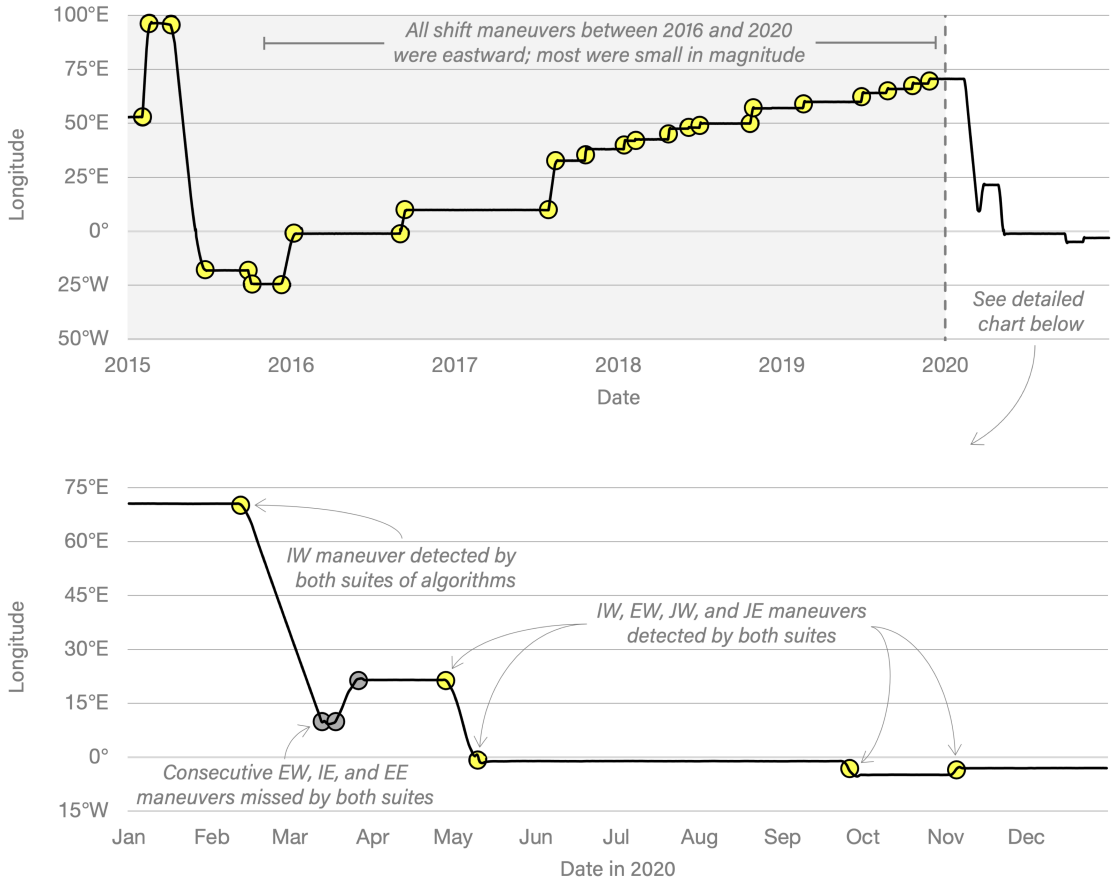

Fig. 14: Maneuver detection results for Russia's *Luch (Olymp)* with daily longitudinal positions from January 1, 2015, to December 31, 2020. The satellite was launched on September 27, 2014, and has occupied more than 20 longitudinal positions in GEO since then, making it one of the most maneuvered satellites in the SPACECOM catalog.

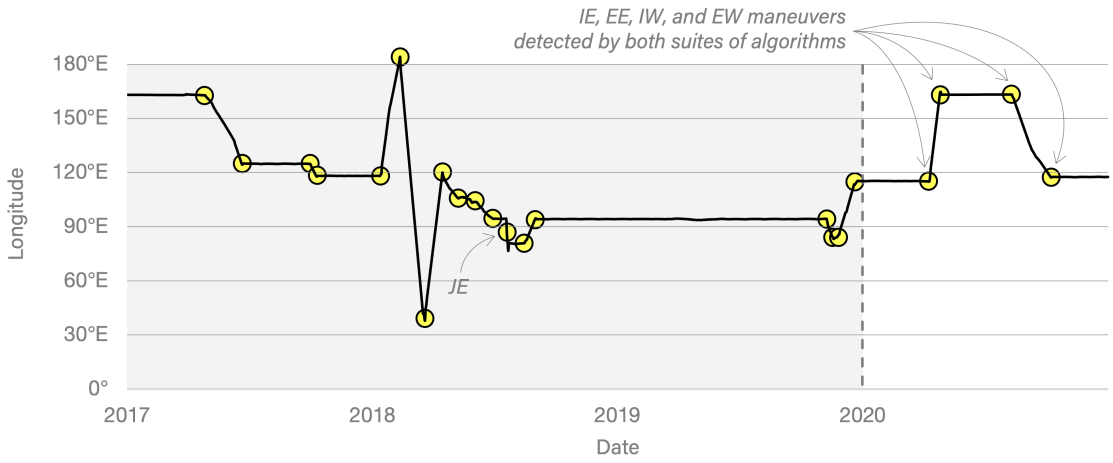

Fig. 15: Maneuver detection results for China's *SJ-17*, including daily longitudinal positions from January 1, 2017, to December 31, 2020. The satellite was launched on November 3, 2016, and has pursued a variety of longitudinal shift maneuvers since then, including several in which the satellite adjusted its drift rate during drift periods.

space community [28]. Unlike other satellites that perform many maneuvers, most of *Luch*'s longitudinal shift maneuvers were small, as shown in the top subplot of Fig. 14. For the majority of *Luch*'s on-orbit lifetime, the satellite has shifted eastward in the geostationary belt, forming a stair-step pattern in its longitudinal position history from the beginning of 2016 to the end of 2019. In 2020, however, *Luch* abandoned this unique PoL and performed two westward longitudinal shifts with relatively long drift periods, among others. The bottom subplot of Fig. 14 describes *Luch*'s behavior during the year 2020. As the figure shows, the two suites of maneuver detection algorithms both detected a majority of the maneuvers performed by *Luch* in 2020, including the initial IW maneuver that indicated that the satellite was deviating from its previous PoL, but failed to correctly identify a cluster of maneuvers (EW, IE, and EE) that were labeled consecutively. Since this cluster of maneuvers is exceedingly rare—it never appeared in the 5- or 10-year training dataset under the labeling strategy described in Subsec. 2.7—the algorithms could not detect it.

SJ-17 is a Chinese GEO satellite that was launched in November 2016. Like *Luch*, *SJ-17* has performed several close approaches during its time in orbit [29]. Unlike *Luch*, many of the satellites approached by *SJ-17* have been Chinese, and likely cooperative partners in various rendezvous and proximity operations [10]. In addition to its pattern of regular close approaches, *SJ-17*'s behavior is also unique in the way it performs its longitudinal shift maneuvers. The satellite has a pattern of adjusting its drift rate during drift periods, creating additional discontinuities in its longitudinal history plot. Despite this pattern continuing into 2020, both suites of maneuver detection algorithms successfully detected all four of the maneuvers labeled in the 2020 test dataset.

5. POLICY APPLICATIONS

Both the method described in this work and a number of its data byproducts—such as the longitudinal shift maneuver database used for algorithm training and the geographic position histories created for all satellites in GEO—have the potential to contribute to several interesting efforts in international space policymaking.

5.1 Studying Compliance with Guidelines and Agreements

In November 2019, the U.S. Government updated its orbital debris mitigation guidelines, which included revised standard practices for post-mission behavior for GEO satellites [1]. According to the new guidelines, GEO satellites should be removed from the geostationary belt at the end of their operational lifetime such that their perigee remains more than 200 km above geostationary altitude for at least 100 years.[6] Additionally, in an effort to minimize "the risk to other space systems from accidental explosions and associated orbital debris after completion of mission operations," all GEO satellites should deplete their onboard propellant and compressed gases once they reach their disposal orbit.

The method for detecting satellite maneuvers in GEO described in the previous sections could contribute to compliance verification for both post-mission disposal of GEO satellites and onboard propellant depletion. Post-mission disposal can easily be detected by this method, as discussed in Sec. 4. The detection of satellite maneuvers of any kind after a satellite reaches its post-mission disposal orbit could suggest the violation of the standard practice of onboard propellant depletion, since some propellant must be expended in order to perform a maneuver.

Unlike operators in other orbital regimes, satellite operators in GEO must acquire a license to operate in the domain from the International Telecommunications Union (ITU), a specialized agency of the United Nations [30]. Licenses from the ITU—which describe a portion of the geostationary belt, known as an *orbital slot*, and a portion of the radio-frequency spectrum, known as a *frequency band*, in which satellites can operate—ensure satellites in GEO have both the physical and spectral space they need to function without harmful radio-frequency interference from neighboring space systems [26]. When satellites perform longitudinal shift maneuvers with a significant longitudinal displacement, they are almost certainly switching orbital slots in the process, which presents challenges when it comes to adhering to ITU license agreements. Detecting such maneuvers is of critical importance to regulators tasked with studying compliance to the ITU's guidelines and evaluating future license applications from operators with previous violations.

This work can also contribute to the verification of public claims of particularly nefarious satellite behavior, even when it is denied by perpetrators. In September 2018, the French defense minister claimed that a Russian satellite, *Luch (Olymp)*, performed an uncoordinated close approach with a French-Italian military satellite, *Athena-Fidus* (Satellite ID: 39509) [19]. Such an action would likely constitute both a violation of ITU license agreements [23] and represent unfavorable behavior within the international space community or an "act of espionage" according to the defense

[6] Note that this implicit definition of GYO differs from the one used in Sec. 1, which was 300 km above geostationary altitude.

minister. Detecting longitudinal shift maneuvers, as demonstrated for this particular Russian satellite in Fig. 14, could help verify France's claim and form the basis of new norms that would prevent future uncoordinated close approaches.

The history of compliance with both the United States' and ITU's guidelines, which can be studied using the geographic position time-histories created for GEO satellites as part of this study, can help inform the development of future space sustainability practices and orbital allocation processes both in GEO and other orbital regimes.

5.2 Developing Future Norms of Behavior

Perhaps the most important contribution of this work to the international space policy community is its potential to form a quantitative basis for the development of future norms of behavior in GEO, including guidance on best practices for longitudinal shift maneuvers and a minimum separation distance between noncooperative satellites [30].

The products of this work—an algorithmic method for detecting longitudinal shift maneuvers and a record of previous maneuvers in an easily interpretable coordinate system—could help form the basis of guidelines surrounding how to perform longitudinal shift maneuvers without interfering with neighboring satellites' activities. The history of GEO satellite operations described in each satellite's geographic position history provides hundreds of examples of how operators choose to pursue maneuvers in the absence of agreed-upon rules and guidelines, which could inform the design of new ones. Additionally, detecting longitudinal shift maneuvers could trigger the investigation of potentially uncoordinated close approaches, which threaten satellites with on-orbit collision and other, non-kinetic types of attacks. A future norm to prevent close approaches, such as setting a minimum separation distance for noncooperative satellites, could be both informed and verified via satellite maneuver detection.

6. CONCLUSION

The results presented in Sec. 3 and discussed in Sec. 4 demonstrate the promise of using supervised machine learning techniques to detect longitudinal shift maneuvers in GEO. Since TLE data is a publicly available resource, this work can be replicated regardless of a researcher's access to other, likely more proprietary space object orbital data, which can make it more easily applicable to policy debates, like those discussed in Sec. 5.

Because this method depends on manually labeling satellite behavior based on geographic position time-histories, it is more limited than some of the other methods cited in Subsec. 1.4. Although satellite behavior can easily be discerned by inspecting historical geographic coordinates for objects in GEO, as shown in Fig. 2, discerning behavior in other orbital regimes in this way is more challenging. In LEO and MEO, for example, satellite positions in latitude and longitude space is rapidly changing over time; inspecting the time-histories of other orbital parameters will likely yield more accurate results. That being said, the dependence on geographic coordinates could continue to be a strength for detecting other interesting behavior in GEO besides longitudinal shift maneuvers, such as identifying and predicting station-keeping maneuvers or classifying satellites' PoLs. Regardless of the types of behavior being detected with this methodology, the algorithm designed in Subsec. 2.7 would likely benefit from hyperparameter tuning, which may lead to more favorable performance metrics than those presented in Sec. 3.

This work both develops and demonstrates a novel method for detecting satellite behavior in GEO, with great potential for other applications, both in on-orbit satellite behavior analysis and developments in international space policy.

7. REFERENCES

[1] U.S. Government Orbital Debris Mitigation Standard Practices, November 2019 Update. Technical report, 2019.

[2] X. Bai, C. Liao, X. Pan, and M. Xu. Mining Two-Line Element Data to Detect Orbital Maneuver for Satellite. *IEEE Access*, pages 129537–129550, 2019.

[3] P. Branco, L. Torgo, and R.P. Ribeiro. A Survey of Predictive Modeling on Imbalanced Domains. *ACM Computing Surveys*, 49(31):31–50, 2016.

[4] N.V. Chawla, K.W. Bowyer, L.O. Hall, and W.P. Kegelmeyer. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.

[5] H.D. Curtis. *Orbital mechanics for engineering students*. Elsevier Aerospace Engineering Series. Elsevier Butterworth Heinemann, 2005.

[6] J. Decoto and P. Loerch. Technique for GEO RSO Station Keeping Characterization and Maneuver Detection. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2015.

[7] P. DiBona, J. Foster, A. Falcone, and M. Czajkowski. Machine learning for RSO maneuver classification and orbital pattern prediction. *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2019.
[8] ESA Space Debris Office. Classification of Geosynchronous Objects. Technical Report GEN-DB-LOG-00211-OPS-GR, European Space Agency, Darmstadt, Germany, 2018.
[9] A. Gulli and S. Pal. *Deep Learning with Keras*. 2017.
[10] T. Harrison, K. Johnson, T.G. Roberts, T. Way, and M. Young. Space Threat Assessment 2020. *Center for Strategic and International Studies*, 2020.
[11] H. He, Y. Bai, E.A. Garcia, and S. Li. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *IEEE International Joint Conference on Neural Networks*, Hong Kong, China, 2008.
[12] H. He and Y. Ma. *Imbalanced Learning: Foundations, Algorithms, and Applications*. Wiley-IEEE Press, Piscataway, New Jersey, 1st edition, 2013.
[13] C. Henry. Intelsat-29e satellite suffers fuel leak, spotted drifting along GEO arc, 2019.
[14] C. Henry. DirecTV fears explosion risk from satellite with damaged battery, 2020.
[15] Inter-Agency Space Debris Coordination Committee Steering Group and Working Group 4. IADC Space Debris Mitigation Guidelines. Technical report, 2007.
[16] A.B. Jenkin, M.E. Sorge, and J.P. McVey. National and International Disposal Requirements and Guidelines Applicable to GPS, 2015.
[17] T. Kelecy, D. Hall, K. Hamada, and D. Stocker. Satellite Maneuver Detection Using Two-line Element (TLE) Data. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, Hawaii, 2007.
[18] M. Kuhn and K. Johnson. *Applied Predictive Modeling*. Springer, New York, New York, 1st edition, 2013.
[19] J. Leicester and S. Corbet. 'Espionage:' French defense head charges Russia of dangerous games in space. *DefenseNews*, 2018.
[20] S. Lemmens and H. Krag. Two-Line-Elements-Based Maneuver Detection Methods for Satellites in Low Earth Orbit. *Journal of Guidance, Control, and Dynamics*, 37:860–868, 2014.
[21] S. Lewin. Old Weather Satellite Headed to 'Graveyard Orbit', 2017.
[22] T. Li, K. Li, and L. Chen. Maneuver Detection Method Based on Probability Distribution Fitting of the Prediction Error. *Journal of Spacecraft and Rockets*, 56:1114–1120, 2019.
[23] S. Mountin. The Legality and Implications of Intentional Interference with Commercial Communication Satellite Signals. *International Law Studies*, 90:101–197, 2014.
[24] N.V. Patel. Boeing and DirecTV are scrambling to move a satellite before it explode, 2020.
[25] Russell P Patera. Space Event Detection Method. *Journal of Spacecraft and Rockets*, 45(3):554–559, 5 2008.
[26] S. Pinnagoda. Harmful Interference and Infringements of the Radio Regulations, 2015.
[27] Rhodes, B.C. PyEphem: Astronomical Ephemeris for Python, 2011.
[28] T.G. Roberts. Unusual Behavior in GEO: Luch (Olymp-K), 2021.
[29] T.G. Roberts. Unusual Behavior in GEO: SJ-17, 2021.
[30] T.G. Roberts and C. Bullock. A sustainable geostationary space environment requires new norms of behavior. *MIT Science Policy Review*, 1:34–38, 2020.
[31] T.G. Roberts and R. Linares. Satellite Repositioning Maneuver Detection in Geosynchronous Orbit Using Two-line Element (TLE) Data. In *International Astronautical Congress*, 2020.
[32] Mike Safyan, Jonathan Rosenblatt, and George John. Exhibit 1: Request for Special Temporary Authority, 2020.
[33] C. Shabarekh, J. Kent-Bryant, G. Keselman, and A. Mitidis. A Novel Method for Satellite Maneuver Prediction. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, Hawaii, 2016.
[34] E.M. Soop. *Handbook of Geostationary Orbits*. Doredrecht, Netherlands, 1994.
[35] U.S. Space Force 18th Space Control Squadron. Space-Track.org.
[36] D.A. Vallado. *Fund. of Astrodynamics and Applications*. Microcosm, Hawthrone, California, 4th edition, 2013.
[37] D.A. Vallado, P. Crawford, R. Hujsak, and T.S. Kelso. Revisiting Spacetrack Report 3. In *AIAA/AAS Astrodynamics Specialist Conference*, Keystone, Colorado, 2006.
[38] S. Zhao and J. Zhang. Minimum-fuel station-change for geostationary satellites using low-thrust considering perturbations. *Acta Astronautica*, 127:296–307, 2016.