End-to-End Behavioral Mode Clustering for Geosynchronous Satellites

Thomas G. Roberts[*]
Massachusetts Institute of Technology
Victor Rodriguez-Fernandez[†]
Universidad Politécnica de Madrid
Peng Mun Siew[‡], **Haley E. Solera**[§], **Richard Linares**[**]
Massachusetts Institute of Technology

ABSTRACT

Unsupervised machine learning techniques can be used to describe historical satellite behavior in the geosynchronous orbital regime with higher fidelity than the colloquial terms most often used by satellite operators and the space situational awareness research community. Although satellite behavior is often described with a handful of shorthand operational modes—east-west station-keeping, libration orbits, and retirement drift in the graveyard orbit, for example—more detailed behavioral modes that differentiate station-keeping maneuver frequencies, magnitudes, and propulsion type can be identified via machine learning. This paper contributes to the GEO satellite pattern-of-life characterization literature in a novel way: dissecting historical GEO satellite position time-histories into short segments and sorting them into groups of similar features using time-series k-means clustering. Once sorted, the populations within each cluster can be analyzed further to uncover similarities that may be differentiating observed behavior, including the type of propulsion mechanisms on board and the age of the satellite bus. Unlike previous efforts in behavioral mode clustering, the algorithm presented in this work can be deployed on long-duration satellite position time-histories across many different behavioral modes, with a single learning step from input to output ("end-to-end"). Additionally, there is no requirement to first identify pattern-of-life (PoL) nodes, in which on-orbit behavior drastically shifts from one behavioral mode to another; instead the presented algorithm can identify segments that contain PoL nodes, effectively identifying the nodes themselves as part of the process.

1. INTRODUCTION

Satellites in geosynchronous Earth orbit (GEO) typically spend the majority of their operational lifetimes performing station-keeping maneuvers to maintain a near-fixed position in the Earth-centered, Earth-fixed inertial reference frame. Although almost all satellite operators command their GEO satellites to perform some kind of station-keeping maneuvers, there is remarkable diversity amongst station-keeping protocols across the regime's population. Some satellites perform infrequent, but high-magnitude maneuvers using chemical thrusters, while others exercise more frequent, long-duration burns using electric propulsion to accomplish the same task [12]. Despite decades of heritage operations in the GEO environment, there is no comprehensive rulebook for GEO satellite behavior. Due to the availability of historical GEO satellite orbital element data and the periodic nature of both perturbative and artificial forces in the geosynchronous space environment, machine learning techniques are particularly well suited to identify and describe patterns of behavior amongst the GEO satellite population, often with a higher fidelity than terms used most frequently in the satellite operator community.

This work uses the GEO satellite patterns of life (PoL) formulation—where satellite behavior can be described as periods of behavioral modes, separated by nodes during which behavior switches from one mode to another—to identify and characterize GEO satellite behavior in a single, end-to-end algorithm via unsupervised machine learning [15]. Using historical geographic longitude positions derived from publicly available two-line element data (TLEs),

[*]Ph.D. Candidate, Department of Aeronautics and Astronautics. E-mail: thomasgr@mit.edu
[†]Associate Professor, Department of Computer Systems Engineering. E-mail: victor.rfernandez@upm.es
[‡]Research Scientist, Department of Aeronautics and Astronautics. E-mail: siewpm@mit.edu
[§]Ph.D. Candidate, Department of Aeronautics and Astronautics. E-mail: hsolera@mit.edu
[**]Associate Professor, Department of Aeronautics and Astronautics. E-mail: linaresr@mit.edu

satellite behavior between PoL nodes can be sorted into groups by a k-means clustering algorithm adapted for time-series data using dynamic time warping (DTW), as shown previously by the author team [13]. The background section offers further insight into how previous methodologies were adapted for this work.

The methodology section describes the application of a time-series k-means clustering algorithm for this problem, including how orbital element data is prepared to be used as input. The results section visualizes the clustered outputs for a particular case study—segments of GEO satellite position time-histories sampled from the entire GEO satellite population from 2010 to 2021—as a two-dimensional map, where objects that appear closer in the visualization are also closer together in the DTW-informed distance formulation embedded in the clustering algorithm. The discussion section interprets the results by commenting on operational similarities between the identified clusters, such as the types of propulsion used by the satellites assigned to each cluster and their average age, and formulates an additional clustering study for further analysis. The concluding section summarizes the work, describes how it can be used for an array of space situational awareness research initiatives, and comments on opportunities for future study.

2. BACKGROUND

This work directly references the PoL framework for describing long-term satellite behavior in GEO, in which satellite activities on orbit can be described as a sequence of behavioral modes—such as station-keeping to one position in the geostationary belt or drifting in longitudinal space either above or below geostationary altitude—separated by nodes [5]. PoL nodes can refer to initiating or ending a longitudinal drift or adjusting the rate of a longitudinal drift once one is already underway, as well as marking the beginning and end of the study period for which GEO satellite position time-histories are being considered by the algorithm.

Previous work by the author team discusses methodologies with which to identify PoL nodes, including by inspection and algorithmically [10, 11, 15]. Since PoL nodes and behavioral modes are perfectly interspersed by definition—a PoL node always separates two behavioral modes and a behavioral mode always starts and ends with a PoL node—identifying PoL nodes effectively identifies behavioral modes, too. Previous work segmented historical GEO satellite position time-histories using node-labeled truth data, ensuring that each segment to be subjected to the clustering algorithm did not itself contain a PoL node, and thus was likely to represent a single behavioral mode. In the methodology described in the following section, GEO satellite position time-histories are considered in their entirety, with no regard to the locations of various PoL nodes, meaning some segments contain one or more PoL nodes, as shown in Figure 1.

The visualizations and analysis strategy presented in the results and discussion sections build on the work of one member of the author team, dedicated to using self-supervised machine learning methods to enhance a variety of space traffic management tasks [16]. The previous work extracted representations of clustering results using Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction, while those presented in the results section of this work use a t-distributed stochastic neighbor embedding (t-SNE) formulation [4, 19].

3. METHODOLOGY

The following subsections describe the data selection and cleaning process, the segmentation strategy with which GEO satellite position time-histories are dissected into smaller pieces, and the clustering algorithm applied to this problem.

3.1 Data Selection and Cleaning

The clustering algorithm takes segments of data from GEO satellite position time-histories, which describe the historical longitudinal positions of a satellite at evenly spaced time intervals. These time-histories can be derived from TLEs published by the U.S. Space Force's 18th Space Defense Squadron on Space-Track.org [18]. This process is documented in previous works by the author team and replicable using a publicly available code repository [5].

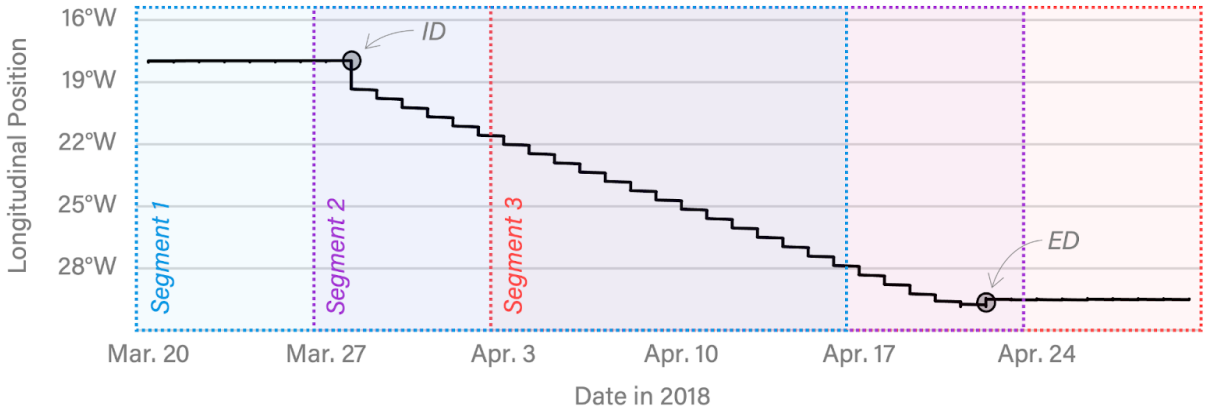

Fig. 1: Segments created with a segmentation strategy with a 28-day kernel size and a 7-day stride. This figure illustrates how six weeks of a GEO position time-history—in this case, describing *Intelsat 901* (Satellite ID: 26824) from March 20 to May 1, 2018—might be broken into segments to feed into a time-series k-means clustering algorithm. During this time, two PoL nodes can be identified by inspection: one corresponding to initiating a westward longitudinal drift ("ID") and another for ending it ("ED"). Since segments cannot be shorter than their defined kernel size, this short GEO position time-history can be divided into just three segments: the first with one enclosed node (shown in blue), the second with two (shown in purple), and the third with one (shown in red).

GEO satellite position time-histories referenced in this study start and end at the bounds of the selected study period—January 1, 2010, to December 31, 2021—and feature a two-hour time step. Position time-histories were created for 918 satellites with at least one TLE on Space-Track.org with an epoch in the bounds of the study period and orbital elements that place it in residence in the protected GEO region as defined by the Interagency Debris Coordination Committee [2]. Not all position time-histories are of equal length: those satellites launched during the study period will not have longitudinal position data for timesteps corresponding to dates prior to their GEO orbital insertion. Large gaps of more than 28 days between TLE epochs for a particular satellite will also lead to gaps in position time-histories, as the orbital elements from TLEs are not propagated more than two weeks forwards or backwards when position time-histories are created.

3.2 Segmentation Strategy

A GEO satellite's longitudinal position at a given moment in time is insufficient to determine the satellite's broader on-orbit behavior. Instead, that position should be contextualized within the positions that come before and after it; these stretches of longitudinal positions can form identifiable signatures of behavior. In this work, that shortlist of longitudinal positions before and after a point of interest is called a *segment*, as it represents a mere portion of the satellites' often-much-longer operational lifetime. Segments are created from position time-histories using a two-parameter segmentation strategy [9]. The first parameter describes the width of each segment: the number of days' worth of longitudinal position data enclosed in each one. In the design of convolutional neural networks (CNNs), this parameter is known as the kernel size. The second parameter describes the displacement between the first timestep enclosed in two adjacent segments, measured in days. In CNNs, this parameter is known as the stride. Figure 1 shows how segments would be created using a segmentation strategy with a 28-day kernel size and a 7-day stride.

Larger kernel sizes allow the clustering algorithm to search for longer-duration patterns, like those associated with sustained station-keeping, but increases the frequency with which PoL nodes are enclosed in segments. When the stride is smaller than the kernel size, individual data points from the longitudinal position time-histories appear in more than one segment, offering multiple opportunities to identify and later characterize the behavior associated with that time step. For the segmentation strategy described in Figure 1, each segment contains 336 longitudinal positions: 12 per day in the segment due to the selected GEO position time-history's two-hour timestep and 28 days in the segment. If a segment is missing one or more internal longitudinal positions, those missing positions are interpolated using a quadratic polynomial interpolation scheme. If a segment is missing one or more longitudinal positions at its beginning or end, the segment is removed from the study and not included in the clustering algorithm described in the following subsection.

All segments created from the 918 GEO satellites included in this study can be fed into the clustering algorithm described in the following subsection or they can be downsampled to decrease the computational resources needed to execute the algorithm. Reasonable strategies for downsampling to consider might include reducing segments to those associated with a shortlist of particular objects of interest or satellites of a particular age. For this study, segments are rebalanced such that three subgroups were approximately equally represented: one with those likely to be describing station-keeping behavior, one with those likely to be describing longitudinal drift behavior, and one with segments that contain one or more PoL nodes. This rebalancing strategy requires a labeled PoL node dataset for the GEO satellite population being studied [13, 15].

As a final pre-processing step, segments should be normalized. The Python package used for time-series clustering in the following subsection, `tslearn`, includes a rescaling tool, which can be used for normalization [17]. Normalizing the segments encourages the clustering algorithm to group together behaviors with similar patterns independent of magnitude. Under this formulation, longitudinal drifts with vastly different drift rates, but the same direction, are more likely to be clustered together. Similarly, station-keeping patterns with similar protocols, but different central longitudinal positions, are also more likely to be clustered together.

3.3 Clustering Algorithm

Because behavioral mode classification requires an understanding of how longitudinal positions change over the course of a segment, it is critical to select a distance metric to feed into the clustering algorithm that accounts for the time-series nature of the data it receives as input. The `tslearn` Python package is equipped with two variations of dynamic time warping (DTW), a tool specifically developed to measure similarities between time series: the DTW Barycenter Averaging (DBA) and soft-DTW algorithms [1, 7, 17]. When these two variants were tested on pre-labeled segment data from GEO satellite position time-histories created using the same methodology as the one described in the previous subsection, the DBA variant produced the smaller sum of the squared error, and was thus selected for this analysis.

The DBA-enabled k-means clustering algorithm from `tslearn` can sort segments into any number of clusters. For this study, the segments were sorted into 24 clusters—a number hypothesized to exceed the number of meaningfully unique behavioral modes in the GEO satellite population—with 10 iterations for the barycenter computational process.

4. RESULTS

The results presented in this section were developed using a segmentation strategy with a 28-day kernel size and a 2-day stride, which produced nearly 1.6 million individual segments. Of these segments, 30,000 were selected for clustering using the downsampling strategy and clustering algorithm described in the previous section. The clustering algorithm was executed using the MIT Lincoln Laboratory's SuperCloud supercomputing system [8].

The clustering results can be visualized in multiple ways. One method, shown in Figure 2, depicts the input segments as thin gray lines in the subplot corresponding to their assigned cluster. Recall that input segments have been normalized, and thus the vertical axis of each subplot in Figure 2 does not directly refer to longitudinal degrees. The centroid of each cluster is shown with a line colored to match one of three broad classes—station-keeping, longitudinal drift, or transitioning between the two. These classes can be labeled by inspection or by noting the number and distribution of PoL nodes contained in each cluster. Henceforth, station-keeping behaviors will be shown in the color blue, longitudinal drift in red, and transitions in purple.

In some cases, like those clusters represented in the first row of Figure 2, classification by inspection is trivial: the centroid of Cluster 1 clearly depicts the signature of a satellite in more than one behavioral mode; Clusters 2 and 4 depict eastward and westward longitudinal drift, respectively; and Cluster 3 depicts station-keeping. In other cases, like determining whether Clusters 5 and 6 represent station-keeping or a transition between two behavioral modes, it may be beneficial to compare the population of each cluster with the number of PoL nodes contained in its segments, should that labeled data be available. Figure 3 shows the population of each of the clusters shown in Figure 2 and the number of PoL nodes contained in each one. Note that the number of PoL nodes in some clusters is higher than that

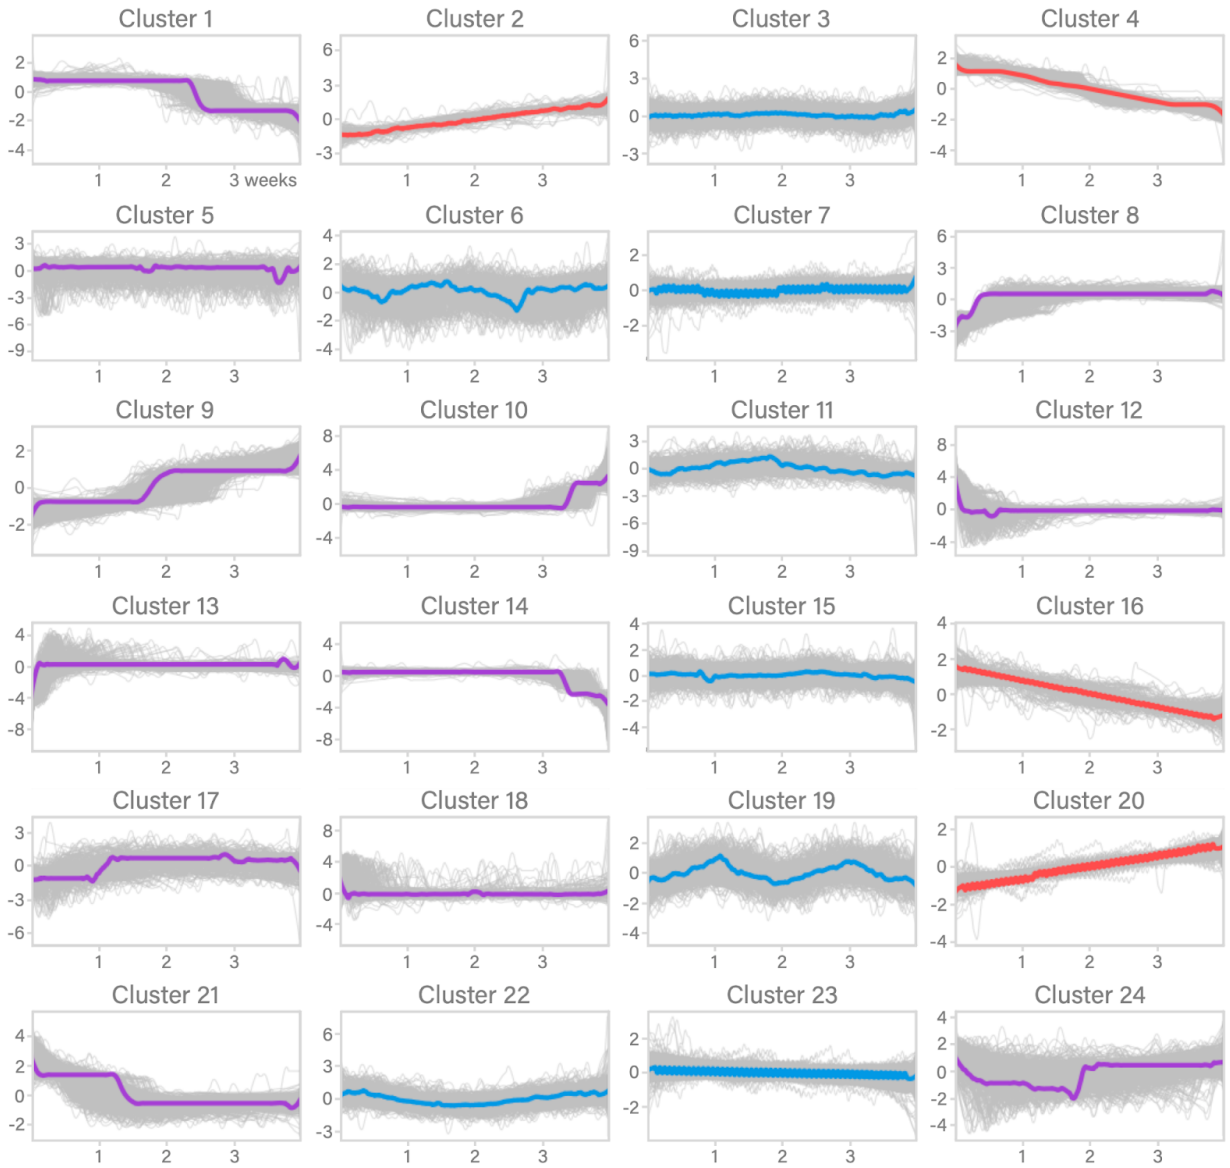

Fig. 2: Segments sorted into 24 clusters. Individual segments—each 28 days long, with normalized values that no longer directly measure longitudinal position in absolute terms—are shown in gray in the background of their corresponding cluster. The centroid of each cluster is shown in the foreground in color. Classes are labeled by a combination of inspection and an analysis of PoL node populations in each cluster. Centroids corresponding to the station-keeping class of behaviors are shown in blue, those corresponding to the longitudinal drift class are in red, and those that correspond to the transition between the two are shown in purple.

cluster's population, because individual segments can, and often do, contain more than one PoL node, as was the case for Segment 2 in Figure 1. Because the number of PoL nodes in Cluster 5 relative to that cluster's population is higher than those clusters that more clearly depict station-keeping behavior, it can be labeled as belonging to the transitions class and colored in purple. Similarly, since the number of PoL nodes in Cluster 6 is small compared to that cluster's population, and the plotted centroid if Figure 2 does not align with the expected behavior of longitudinal drift, it can be labeled as belonging to the station-keeping class and colored in blue. A similar methodology of combining classification by inspection and population analysis can be applied to the remaining clusters.

Although Figure 2 works well to compare segment behavior across clusters, it fails to describe how similar the various clusters are to one another. Each segment appears just once in the visualization, with no indication about whether

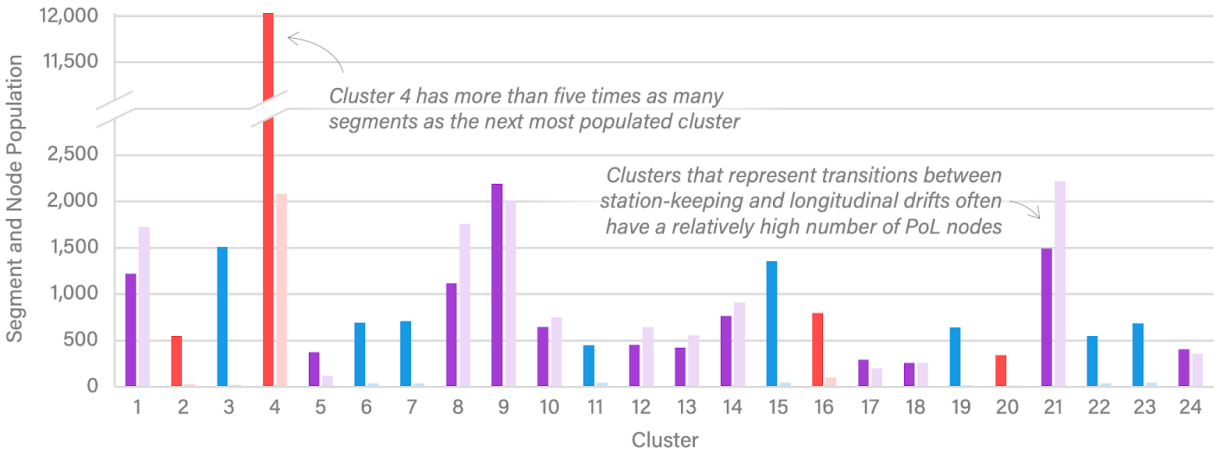

Fig. 3: Segment and PoL node population of each cluster. Segment populations are shown in dark colors; PoL node populations are shown in light colors. The number of PoL nodes relative to the segment population of a cluster can be used to sort the clusters into the three broad behavioral classes: station-keeping, shown in blue; longitudinal shift, shown in red; and transitioning between the two, shown in purple.

segments were near or distant to other identified clusters in the DBA distance space. Figure 4 sheds light on these patterns by representing each segment as a single point on a dimension-reduced representation of the cluster space, where points are closer to some clusters' centroids than they are to others, based on the distance results outputted by the algorithm. Figure 4's three subplots refer to the broad behavioral classes identified in Figure 2, each created using t-SNE for dimensionality reduction with a perplexity value of 100 [19].

Note that the horizontal and vertical dimensions of Figure 4 hold no significance in the physical GEO environment.

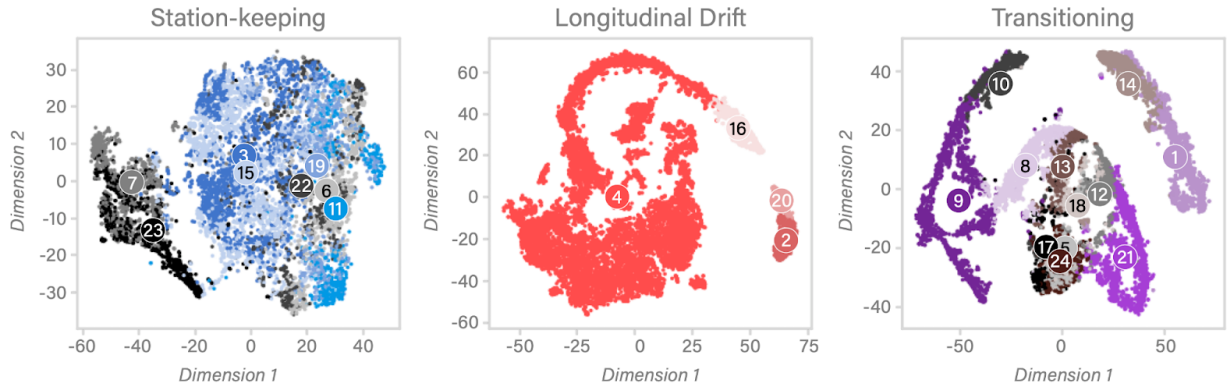

Fig. 4: Two-dimensional representations of the cluster space by behavioral class. In this visualization, created using t-SNE, each point represents a segment. Like Figure 2, each segment appears just once in the plot above, in the subplot corresponding to their broader class: station-keeping at left, longitudinal drift in the center, and transitions at right. Each of the 24 clusters' centroids are labeled with their corresponding cluster index.

5. DISCUSSION

The relative distances between cluster centroids in Figure 4 match intuitive readings from Figure 2. In the station-keeping class, the centroids for Clusters 7 and 23 are relatively close in Figure 4, and their corresponding segments are distant from those in other clusters in the same class. Unsurprisingly, this same pair of clusters exhibit similar patterns in Figure 2: oscillations around a near constant longitudinal position. Clusters 3 and 15 are also adjacent in Figure 4 and similar in Figure 2: they appear to exhibit the flattest centroids of the remaining station-keeping clusters.

Patterns in the longitudinal drift class are even easier to assess. Cluster 4—by far the largest class, as shown in Figure 3—is adjacent to Cluster 16 in Figure 4. The two clusters clearly correspond to westward drift when inspected in Figure 2. Clusters 2 and 20 are adjacent in Figure 4 and separated from the others. These smaller clusters depict eastward drift in Figure 2, the less common of the two drift directions in GEO. Similar analysis could be pursued on the third subplot of Figure 4 to better understand how PoL nodes are recognized by the clustering algorithm, but that is not the focus of this study.

Beyond the similarity of the longitudinal position signatures in Figure 2, what satellite characteristics could shed light on the distribution of clusters in Figure 4? Using data from Slingshot Aerospace's Seradata space object catalog, the segments in each cluster can be assigned propulsion types: chemical, electric, hybrid, or unknown [14]. The start date of each segment can also be compared to the corresponding satellite's launch date in the Space-Track.org catalog to provide summary statistics on the age of the satellite buses in each cluster.

On average, segments in both Clusters 7 and 23 correspond to satellites that have been in GEO longer than average—21.4 and 23.6 years, respectively—longer than those sorted into the other station-keeping clusters. Cluster 3 and 15 correspond to satellites of a more modest age—approximately nine years—while the remaining four clusters all correspond to younger satellites. The segments sorted into Cluster 11 disproportionately correspond to satellites equipped with electric propulsion. The same is true for Cluster 11's adjacent neighbors in Figure 4—Clusters 6, 19

Table 1: Average characteristics for segments' corresponding satellites by cluster.

Cluster	Behavioral Class	Propulsion Type				Satellite Age (years)	
		Chemical	Electric	Hybrid	Unknown	Mean	Standard Deviation
1	Transitioning	63.5%	1.1%	21.3%	14.2%	11.8	7.4
2	Longitudinal drift	26.5%	2.0%	9.8%	61.6%	22.4	11.4
3	Station-keeping	71.2%	0.3%	20.4%	8.1%	9.4	6.0
4	Longitudinal drift	68.8%	2.0%	10.5%	18.7%	19.8	7.8
5	Transitioning	41.6%	2.4%	44.0%	11.9%	5.6	4.9
6	Station-keeping	59.9%	3.5%	22.2%	14.4%	6.8	4.8
7	Station-keeping	69.5%	0.1%	11.3%	19.1%	21.4	9.5
8	Transitioning	55.3%	2.0%	31.3%	11.5%	6.5	7.2
9	Transitioning	59.5%	2.2%	18.7%	19.6%	11.7	8.8
10	Transitioning	57.9%	1.1%	28.4%	12.6%	9.3	7.2
11	Station-keeping	38.7%	8.6%	33.8%	18.8%	6.9	4.9
12	Transitioning	53.7%	1.1%	24.1%	21.1%	5.3	7.1
13	Transitioning	55.8%	1.4%	29.2%	13.6%	5.1	6.7
14	Transitioning	66.4%	0.6%	19.5%	13.5%	12.5	7.6
15	Station-keeping	66.2%	1.5%	18.5%	13.8%	8.9	5.7
16	Longitudinal drift	29.4%	1.8%	10.8%	58.0%	23.1	10.8
17	Transitioning	37.7%	5.1%	38.0%	19.2%	5.2	5.3
18	Transitioning	48.8%	2.3%	32.3%	16.5%	5.3	6.6
19	Station-keeping	69.6%	2.3%	19.3%	8.9%	7.8	4.9
20	Longitudinal drift	23.2%	1.2%	15.9%	59.7%	26.8	8.7
21	Transitioning	52.2%	1.4%	24.4%	22.0%	6.1	7.4
22	Station-keeping	59.7%	1.8%	22.2%	16.2%	8.2	5.3
23	Station-keeping	56.0%	0.7%	17.4%	25.8%	23.6	8.8
24	Transitioning	48.0%	2.2%	31.9%	17.9%	4.5	5.9

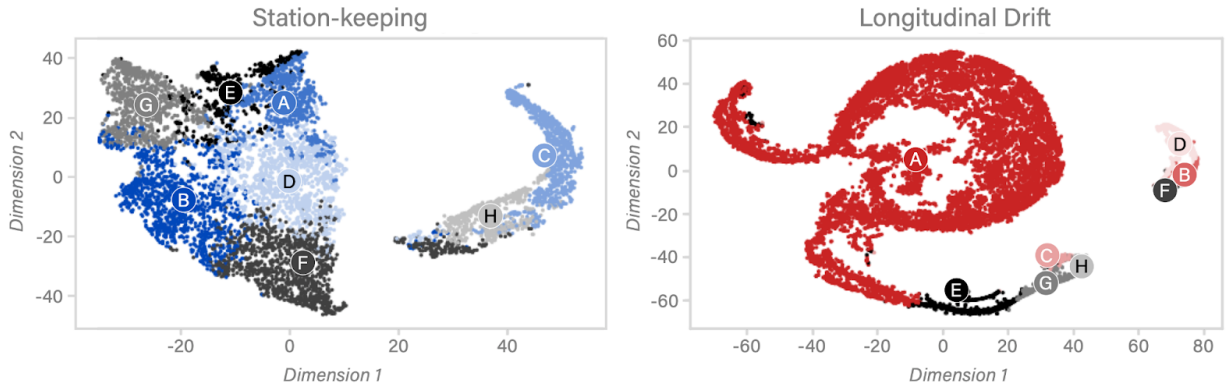

Fig. 5: Two-dimensional representations of the cluster space after first dividing segments by behavioral class. This visualization, created using t-SNE with a perplexity parameter of 100, represents the results of an independent execution of the clustering algorithm. Thus the eight clusters shown in each subplot do not correspond to one another, nor the clusters shown in Figure 4. Instead, they offer an illustration of how one clustering process can inform another.

and 22— but to a lesser extent. In the longitudinal drift class, Cluster 4 is an outlier because of the propulsion types featured on its corresponding satellites: 68.8 percent, more than twice as high a fraction as the other longitudinal drift clusters. Cluster 20, Cluster 4's neighbor in Figure 4, corresponds to more hybrid-propulsion satellites than those with segments sorted into Clusters 2 and 20.

To clarify the differences between the clusters within the station-keeping and longitudinal drift behavioral classes, the original segments can be first divided into the classes assigned to each cluster in Figure 2 and then fed through the same clustering algorithm. Figure 5 shows the dimension-reduced cluster results space for eight clusters in both the station-keeping and longitudinal drift classes. Here, the clusters in the station-keeping class show less overlap than those in Figure 4, but the distinction between new and old satellites still dominates the distribution, suggesting a pattern of observable shifts in station-keeping protocols as satellites age. Clustering the longitudinal shift class into eight clusters—allowing analysts to consider the possibility that there are more meaningfully unique behavioral modes than the four identified in Figure 4—breaks the westward drift group into five clusters and the eastward drift group into three. Further analysis on the available satellite characteristics of these clusters could offer further insights on how the algorithm may have produced the results seen in Figure 5. Alternatively, such analysis could reveal that the algorithm was asked to identify more clusters than there are meaningfully differentiable behavioral modes. In that case, clusters with similar features could be combined and the number of unique modes could be used to inform the number of clusters on future iterations of the algorithm.

Future research should explore concepts of soft clustering (where the probability of belonging to each cluster is computed), test the effect of varying kernel size and stride in the segmentation strategy, and compare the observed results of the k-means clustering algorithm with a hierarchical clustering one, such as CHAMELEON or Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) [3, 6].

6. CONCLUSIONS

One output of this work—segment-by-segment results for every satellite in a study list—can be used to describe a GEO satellite's PoL as a history of cluster assignments. A satellite that is first inserted into the GEO belt, then maintains a single longitudinal position with high-frequency east-west and north-south station-keeping, and finally retires to the graveyard orbit, all within the study period, would perhaps be grouped into only a small number of clusters: those associated with its three behavioral modes and the two nodes during which it was transitioning. Segment-level results could describe the satellites' PoL on a day-by-day or week-by-week basis.

In addition to contributing to the burgeoning literature on satellite pattern-of-life characterization, this work can also be used to describe historical GEO satellite station-keeping behavior at low resolution. That is, a satellite's historical

behavior could be described by behavioral mode classes and short lists of relevant parameters, as opposed to the high-resolution longitudinal position time-history data on which the algorithm was trained.

Together, this sort of information could contribute to international coordination in GEO, specifically the evaluation of the use of orbital allocations from the International Telecommunication Union (ITU), a specialized agency of the United Nations. Future studies could use PoL behavioral mode assignments created from this work to inform compliance studies. That is, when the end-to-end clustering algorithm identifies that a satellite is in one of many station-keeping modes, the corresponding satellite could be assessed for compliance with the physical orbital position described in active ITU satellite networks, and thus protected from harmful interference in the radio-frequency spectrum.

In short, the results from this work demonstrate that machine learning techniques show promise for assisting human-led analysis efforts in space situational awareness: desaturating the high-volume, high-dimensional data associated with historical space object tracking while preserving key insights that describe how satellites in GEO are controlled.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2141064. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing high-performance computing resources that have contributed to the research results reported within this work.

REFERENCES

[1] M. Cuturi and M. Blondel, “Soft-DTW: a differentiable loss function for time-series,” arXiv:1703.01541, 2017.
[2] Inter-Agency Space Debris Coordination Committee Steering Group and Working Group 4, “IADC Space Debris Mitigation Guidelines," 2007.
[3] G. Karypis, Eui-Hong Han, and V. Kumar, “Chameleon: hierarchical clustering using dynamic modeling,” *Computer* 32 (8), 1999.
[4] L. McInnes, J. Healy, and J. Melville, “UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction,” 2020, arXiv:1802.03426.
[5] MIT Astrodynamics, Space Robotics, and Controls Laboratory, “Geosynchronous Satellite Pattern of Life Characterization,” 2023, https://aeroastro.mit.edu/arclab/research/geosynchronous-satellite-pattern-of-life-characterization/.
[6] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research* 12: 2825-2830, 2011, https://scikit-learn.org/stable/modules/generated/sklearn.cluster.HDBSCAN.html.
[7] F. Petitjean, A. Ketterlin, P. Gançarski, “A global averaging method for dynamic time warping, with applications to clustering,” *Pattern Recognition* 44 (3), 2011.
[8] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas, “Interactive Supercomputing on 40,000 Cores for Machine Learning and Data Analysis,” IEEE High Performance Extreme Computing Conference (HPEC), 2018.
[9] T. G. Roberts, “Geosynchronous Satellite Maneuver Classification and Orbital Pattern Anomaly Detection via Supervised Machine Learning,” Master's thesis, Massachusetts Institute of Technology, 2021.
[10] T. G. Roberts and R. Linares, “Geosynchronous Satellite Maneuver Classification via Supervised Machine Learning,” 22nd Advanced Maui Optical and Space Surveillance Technologies Conference, 2021.

[11] T. G. Roberts and R. Linares, "Satellite Repositioning Maneuver Detection in Geosynchronous Orbit Using Two-line Element (TLE) Data," 71st International Astronautical Congress, 2020.
[12] T.G. Roberts and R. Linares, "A Survey of Longitudinal-Shift Maneuvers Performed by Geosynchronous Satellites from 2010 to 2021," 73rd International Astronautical Congress, 2022.
[13] T. G. Roberts, H. E. Solera, and R. Linares, "Geosynchronous Satellite Behavior Classification via Unsupervised Machine Learning," 9th International Academy of Astronautics Space Traffic Management Conference, 2023.
[14] Slingshot Aerospace, "Seradata," https://spacetrak.seradata.com/.
[15] H. E. Solera, T. G. Roberts, and R. Linares, "Geosynchronous Satellite Pattern of Life Node Detection and Classification," 9th International Academy of Astronautics Space Traffic Management Conference, 2023.
[16] E. Stevenson, V. Rodriguez-Fernandez, H. Urrutxua, V. Morand, and D. Camacho, "Self-supervised Machine Learning Based Approach to Orbit Modelling Applied to Space Traffic Management," 11th International Association for the Advancement of Space Safety (IAASS) Conference, 2021.
[17] R. Tavenard, J. Faouzi, G. Vandewiele, F. Divo, G. Androz, C. Holtz, M. Payne, R. Yurchak, M. Rußwurm, K. Kolar, and E. Woods, "Tslearn, A Machine Learning Toolkit for Time Series Data," *Journal of Machine Learning Research* 21 (118): 1-6, 2020.
[18] U.S. Space Force 18th Space Defense Squadron, Space-Track.org.
[19] L. van der Maaten and G. Hinton, "Visualizing Data using t-SNE," *Journal of Machine Learning Research* 9 (11), 2008.