Space Object Attitude Stability Determined from Radar Cross-Section Statistics

Matthew A. Stevenson
LeoLabs, Inc
Michael Nicolls and Chris Rosner
LeoLabs, Inc

ABSTRACT

We present a new method for estimating the attitude stability of low-Earth resident space objects (RSOs) using radar cross-section (RCS) statistics from LeoLabs' global radar network. Assuming a non-isotropic shape, an Earth-oriented RSO will have an elevation-angle dependent RCS when viewed from a ground-based radar. We thus test RSO attitude stability by looking for a difference in median RCS when the satellite is viewed at different elevation angles. We demonstrate this technique using data from LeoLabs' Midland Space Radar. This has applications for interpreting RSO stability, and its outputs can be used for advanced object and mission classification.

1. INTRODUCTION

An essential requirement for Space Situational Awareness (SSA) is the characterization of Resident Space Objects (RSOs). An essential characteristic of RSOs is their attitude stability[3]. Specifically, it is of taxonomical interest whether an RSO has a controlled, stable attitude or a tumbling or irregular attitude. Further, tracking changes in the attitude stability of an RSO allows for the detection in on-orbit changes in the characteristics of that RSO, be they due to disabling or enabling of payload systems, RSO fragmentation, or some other cause.

The characterization of RSO attitude stability is challenging for several reasons. Resolved optical imaging of many RSOs is prohibitive due to the physical distance and size of the objects. This has led to a degeneracy in measurements between the attitude properties of an RSO and its physical shape. Studies have been performed to show the viability of separating the attitude and shape characteristics of RSOs using high time resolution photometric or radiometric measurements[4][5]. These approaches have promise, but may be prohibitively expensive when scaled up to the larger RSO population.

There are an estimated 200,000 RSOs in Low-Earth Orbit (LEO) of size greater than 2cm [1]. In order to characterize the attitude stability of all of these objects, a low-cost, scalable method is needed.

We present a new method for estimating the attitude stability of spatially unresolved RSOs in LEO using statistics of radar cross-section (RCS) measurements from LeoLabs' radar facilities. These RCS measurements are produced as a natural byproduct of our orbital tracking measurements, and do not require additional radar resources. This method's sole assumption about the shape of the RSOs is that their RCS is non-isotropic with respect to aspect angle. This assumption tends to hold true for RSOs with unequal spatial dimensions. For example, a satellites with large, deployed solar panels tend to satisfy this assumption, while spherical satellies tend not to.

This method uses the statistics of RCS measurements of an RSO as it passes above our radar over many orbits and a wide variety of aspect angles. Using these RCS statistics, we construct a Stability Index (SI). Section 2 describes the SI, Section 3 explains how we use the SI to classify the attitude stability of RSOs, and Section 4 presents several case studies demonstrating practical applications of the SI. Section 5 discusses how we expect the SI to improve in the future. Finally, Section 6 provides conclusions to this work.

LeoLabs operates and continues to build a network of dedicated radar sites around the globe for the purpose of tracking RSOs in LEO. The work of this paper makes use of the data from the Midland Space Radar (MSR) in Midland, Texas. We anticipate extending this work to our next generation radar facilities in New Zealand and other locations.

2. STABILITY INDEX

We seek to characterize the attitude stability of an RSO via its RCS statistics. In this work, an attitude stabilized RSO is defined as one whose attitude is nearly constant relative to the orbital frame. This work assumes that the RSO has different median RCS values when viewed from different aspect angles.

It is important that we work with median RCS values, rather than individual RCS measurements. This is because small changes in aspect angle can result in large changes in measured RCS for complex radar targets. This effect is so pronounced that it is often convenient to treat RCS measurements as drawn randomly from a probability distribution function[7]. Thus, estimating the median of the measured RCS values over a period of several weeks provides a more direct probe of the target's radar reflection properties than any one individual RCS measurement. We use the median rather than the mean of the RCS measurements to make the result less sensitive to outlier RCS measurements.

If an RSO with non-isotropic RCS is attitude stabilized, then its median RCS will vary with the topocentric view-angle with which it is measured (see Fig. 1). Conversely, if the RSO is tumbling, then it will present a random aspect to the radar regardless of topocentric view-angle. This will lead to a median RCS which is independent of topocentric view angle.

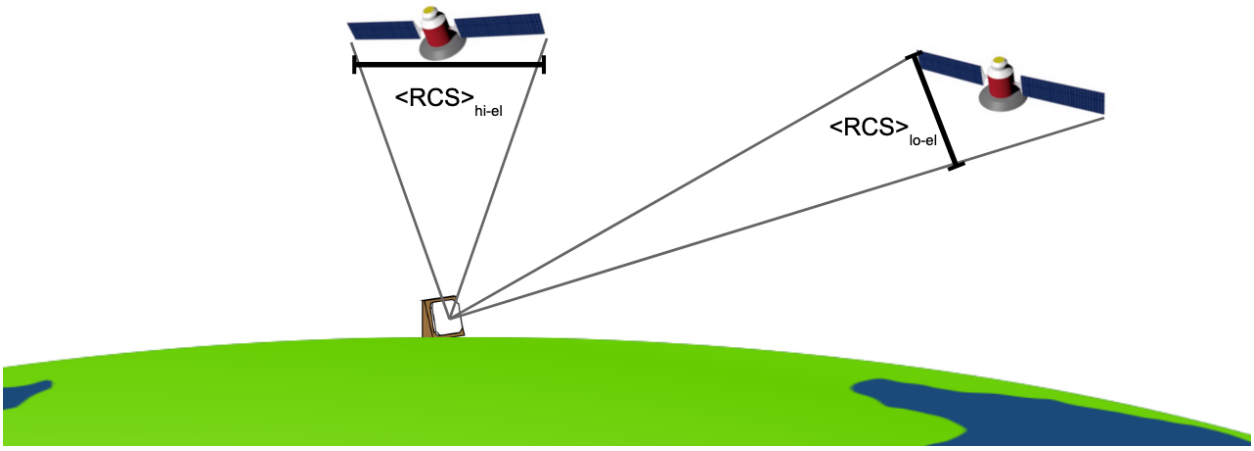

Fig. 1: Schematic drawing of the approximate mean RCS of a stabilized satellite, as viewed from different aspect angles. When cartoon satellite is viewed lower elevation relative to the radar site, it presents a smaller geometric cross-section to the radar, resulting in a lower mean RCS.

Our metric for attitude stabilization exploits this concept by comparing the median RCS of a satellite measured at high elevation to the median RCS of that satellite measured at low elevation. More precisely:

$$\mathrm{SI} = \log_{10} \frac{\langle \mathrm{RCS}_{\mathrm{hi-el}} \rangle}{\langle \mathrm{RCS}_{\mathrm{lo-el}} \rangle} \tag{1}$$

For an attitude-stabilized RSO with non-isotropic RCS, the SI of equation 1 will tend to be non-zero. If the RSO presents its largest RCS to nadir, then the Stability Index will tend to be positive. If the RSO presents its smallesst RCS to nadir, then the Stability Index will tend to be negative.

The use of the logarithm is chosen for two reasons: the ratio is observed to span several orders of magnitude, and the logarithm has symmetric scaling when $\langle \mathrm{RCS}_{\mathrm{hi-el}} \rangle > \langle \mathrm{RCS}_{\mathrm{lo-el}} \rangle$ and when $\langle \mathrm{RCS}_{\mathrm{hi-el}} \rangle < \langle \mathrm{RCS}_{\mathrm{lo-el}} \rangle$.

In our implementation of this metric, we have chosen to use a weighted median over a 60 day timescale. The weighted median was selected because it combines the median's lack of sensitivity to outliers with the weighted mean's emphasis of slowly varying structure over uncorrelated statistical noise. The 60 day timescale was selected because it was observed to provide significant additional suppression of short-term variation in archival MSR data. As we continue to build additional radar facilities, we expect the increased amount of data to allow us to shorten this timescale significantly. This is discussed further in Section 5.

We calculate the weighted median of a set of unevenly time-sampled RCS measurements as follows. When a new RCS measurement for a given RSO is made at high or low elevation, we select all previous RCS measurements for that RSO at high or low elevation in the past 60 days. We then assign a weight to each of those RCS measurements using a Hamming window, with end point at the present measurement and start point 60 days prior. Then the RCS measurements, and their corresponding weights, are sorted in increasing RCS order. The cumulative sum of the sorted weights is calculated, and the index at which that cumulative sum exceeds half of the sum of all weights is index of the weighted median RCS.

After each update to the weighted median RCS for an RSO, whether at high or low elevation, the SI is re-calculated. In this way, we form a time series of the SI, which is updated whenever we make an RCS measurement of the RSO.

We emphasize that the non-isotropic RCS assumption is critical. This approach is not capable of distinguishing an attitude-stabilized sphere from a tumbling sphere, for example. Therefore, a null SI detection can be consistent with an attitude-controlled RSO with an isotropic RCS.

We also note that some attitude-controlled satellites may be commanded to have many different attitudes during a 60-day period. The SI is designed to detect if RSOs always present the same aspect to nadir, so an active satellite that is undergoing desired attitude changes many times during an SI calculation time window will tend to have a null SI value. The 60-day period is expected to decrease as we build additional radar facilities, so we expect to gain SI sensitivity for such satellites in the future.

3. STABILITY CLASSIFICATION

Classifying RSOs based upon the SI is non-trivial. We must balance the competing concerns of designing a classification which is responsive to changes in RSO behavior, while minimizing confusion due to the statistical nature of RCS measurements. This is further complicated by the fact that there is an extremely limited quantity of reference data upon which to train a classification algorithm.

Our approach is to use extensive Monte Carlo simulations with RCS values drawn from an empirically chosen distribution function. We have found that we are able to set a target false-positive rate for stability classification and back-calculate the appropriate Stability Index thresholds as a function of the number of measurements used and their signal-to-noise ratio (SNR). Furthermore, we simulate an ensemble of SI time series in order to calculate a necessary period of time for the SI to remain across a threshold before we change the RSO's stability classification. As with the SI thresholds, these time thresholds are selected to give a target false-positive rate for changes in stability classification.

Our measured RCS values have been found to follow a log-normal distribution.

$$P(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right) \tag{2}$$

The probability distribution function (PDF) of our RCS measurements was estimated by taking our RCS measurements for a given RSO, normalizing those measurements by the median RCS of that RSO, and then estimating the PDF by constructing a histogram. This was repeated for all RSO's for which we have RCS measurements. The resulting PDF estimates are shown in Fig. 2, as well as an analytical log-normal curve with $\mu = 0$ and $\sigma = 1.2$. This is not unexpected, as a log-normal distribution is often used to model RCS measurements[7]. Therefore, our use of a log-normal PDF for RCS Monte Carlo simulations is justified.

The Monte Carlo selection of SI thresholds is performed by simulation a large ensemble of RSOs with various isotropic RCS values and numbers of measurements. For a given iteration, the timing of the measurements are drawn from the Dirichlet distribution (to simulate infrequent RCS measurements) over a 60 day period, and the elevations of those measurements were drawn from a uniform distribution. The actual RCS values are drawn from the log-normal distribution with $\mu = 0$ and $\sigma = 1.2$. Given our known radar properties, we calculate the expected SNR for those measurements, allowing us to add additional Gaussian noise (to simulate measurement noise) and apply a threshold to simulate our radar sensitivity limit [1]. Thus we are able to estimate the SI for this iteration of the Monte Carlo simulation. This process is repeated to create a large ensemble of RSOs with varying RCS values and numbers of measurements.

[1] Note that the RCS-to-SNR conversion is elevation dependent, which introduces a noticeable bias in the SI.

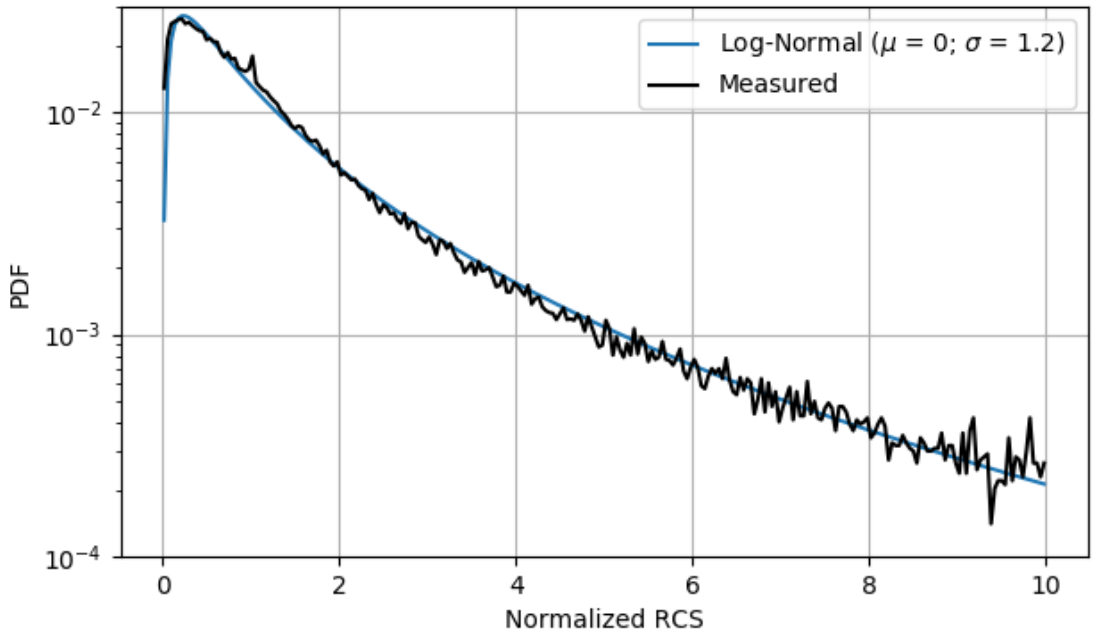

Fig. 2: Measured PDF of RCS measurements for all RSOs, normalized to $\langle RCS \rangle = 1$. A log-normal curve is plotted for reference, with $\mu = 0$ and $\sigma = 1.2$. It is clear that the log-normal curve offers a good description of the measured PDF out to a normalized RCS of 10.

We then place the ensemble of SI values into bin according to their SNR values and statistical weights. The statistical weight is defined as the sum of all Hamming weights of the measurements for the given SI value. In each bin, we construct a histogram of the SI values and select SI thresholds allowing 1% and 10% false positive rates. Finally, we note that these thresholds vary approximately linearly with SNR and numbers of measurements, so we fit a linear function to them. Using this linearly varying threshold function, we are able to calculate the appropriate thresholds for a given SI value given the SNR and number of the individual measurements that led to that SI value.

An example of this for the 10% threshold is shown in Fig. 3. In this figure, there is clear bias towards negative SI values due the the measurement SNR threshold. Perhaps counter-intuitively, the thresholds become higher as SNR increases. The reason for this is that the variance of the log-normal distribution increases with its median. Thus, because SNR is correlated with RCS, the higher SNR targets will tend to have greater variance in their SI metrics.

With SI thresholds in hand, we simulate an ensemble of SI time series and their corresponding 1% and 10% thresholds. We then select time thresholds such that only 10% of isotropic RCS RSOs will cross the SI thresholds for that period of time in a year of measurements. This is shown in Fig. 4. Thus we compute the time thresholds for our stability classification.

The stability classification algorithm applies the above SI and time thresholds to our data via a state machine. Each RSO is initially assumed to be in the non-stabilized state prior to data collection. If, during the course of RCS measurements, that RSO is found to have an SI beyond the thresholds, then the stability classification state of that RSO will remain unchanged. A stability classification state transition will only occur after the SI is seen to be beyond the threshold for a period of time greater than the time thresholds. This algorithm is illustrated in Fig. 5.

Fig. 6 shows example SI curves for two RSOs. The top panel shows Iridium 134 (NORAD ID 43075), which has a clear, consistent stability signal. Our stability classification algorithm identifies Iridium 134 as stabilized. The bottom panel shows Envisat (NORAD ID 27386) with no clear stability signal. Our algorithm classifies Envisat as having no evidence for stabilization. This is consistent with independent measurements showing evidence for the uncontrolled tumbling of Envisat[8].

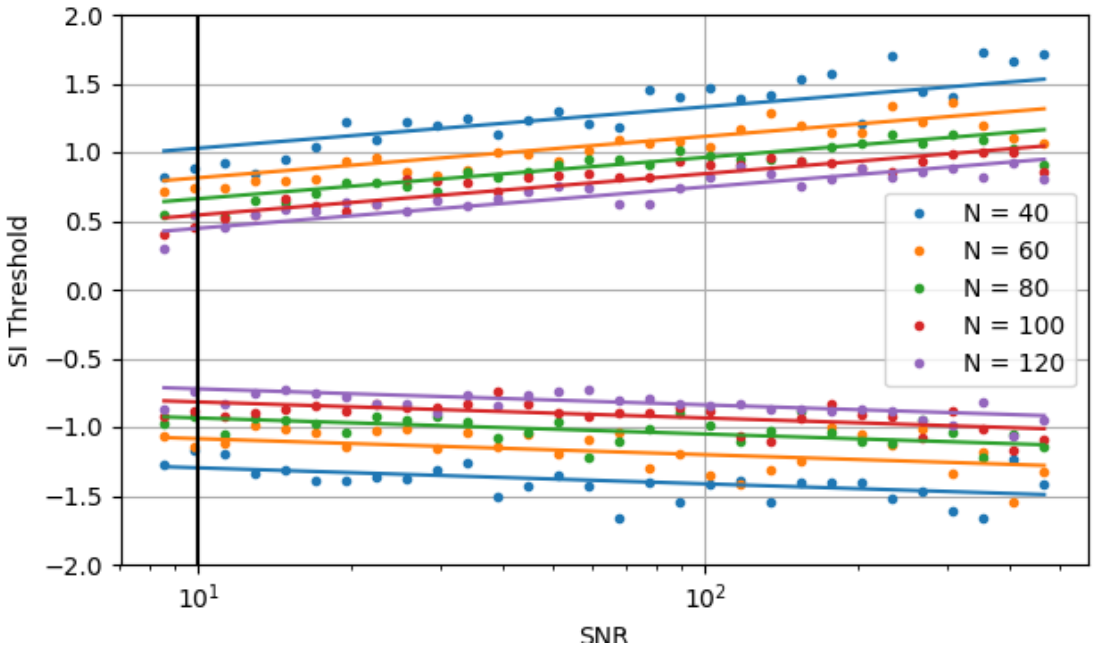

Fig. 3: Monte carlo results for the 10% false-positive rate thresholds as a function of measurement SNR and statistical weight N. The linear model to the thresholds is shown as the solid lines.

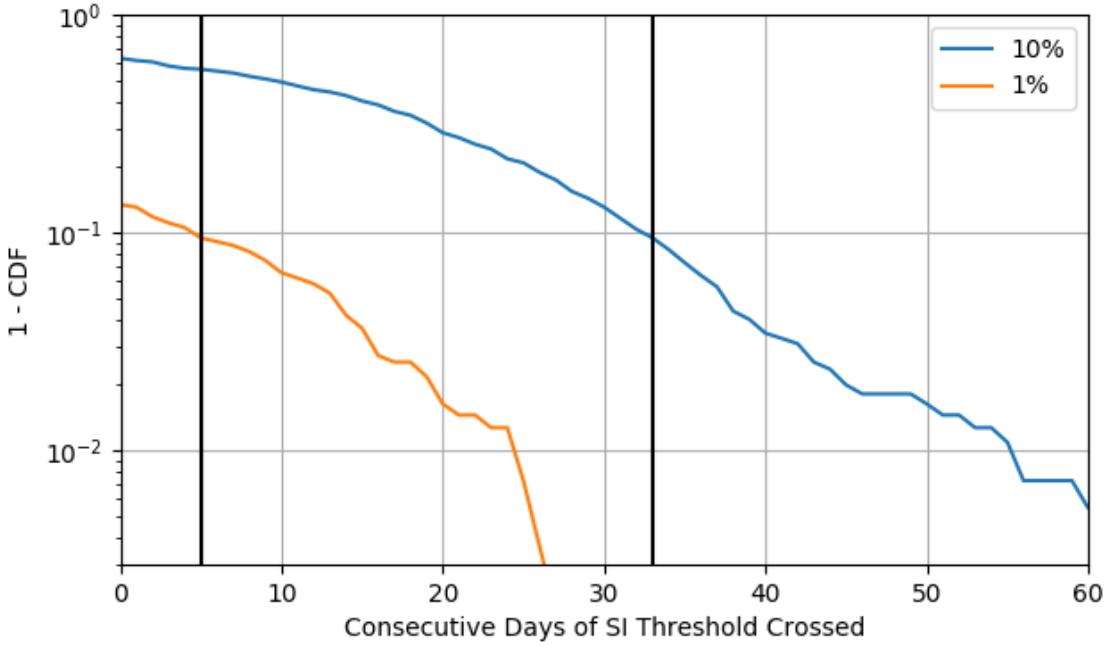

Fig. 4: Monte carlo results for the probabilty than an isotropic RCS RSO will have an SI across the 1% or 10% threshold for a given number of consecutive days in one year. This is calculated as one minus the cumulative distribution function (CDF) of the number of consecutive days that such an RSO is across the SI threshold. The vertical black lines show the time threshold giving a 10% annual false-positive rate.

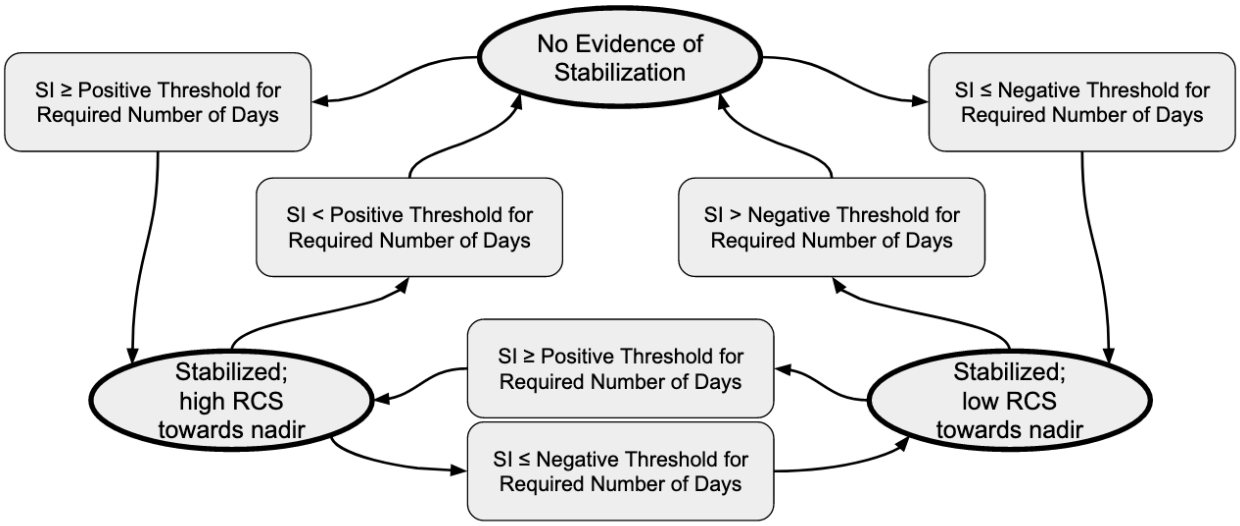

Fig. 5: Stability classification algorithm. If the SI exceeds a positive or negative threshold for the required number of consecutive days, then a stability classification state transition occurs.

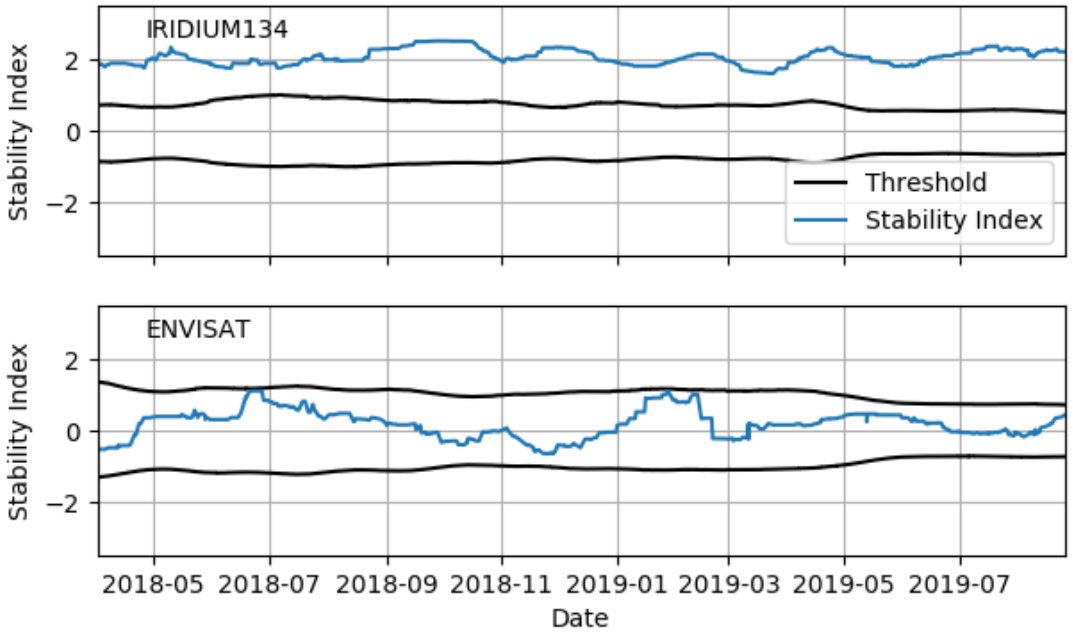

Fig. 6: Stability Index time series for the Iridium 134 (NORAD ID 43075) and Envisat (NORAD ID 27386) satellites. The black curves are the 10% false-positive thresholds for these data sets. Iridium 134 shows a clear, consistent attitude stabilization signal, while Envisat shows no such signal.

4. EXAMPLE APPLICATIONS

4.1 TDS-1 Deorbit Sail Deployment

TechDemoSat-1 (TDS-1, NORAD ID 40076) is an attitude-controlled satellite built by a consortium led by Surrey Satellite Technology Ltd. It contains a wide variety of technology demonstration payloads, including a Deorbit Sail (DOS). The DOS was deployed on April 24, 2019.

The SI for TDS-1 is shown in Fig. 7. The date of the DOS deploy is indicated by a vertical black line. Prior to the deployment, the SI was hovering near the positive threshold. Immediately after the sail deployment, the SI began to show much greater variability. This shows a clear change in the attitude behavior of TDS-1, which is consistent with either a change in RSO attitude (such that the maximum RCS is no longer presented to nadir) or a loss of attitude control.

Fig. 7 also shows the stability classification state for TDS-1 during this period of time. Prior to the DOS deploy event, the SI was near the SI threshold, causing the state to waver. We discuss in Section 5 how the future addition of new radar sites will improve this performance.

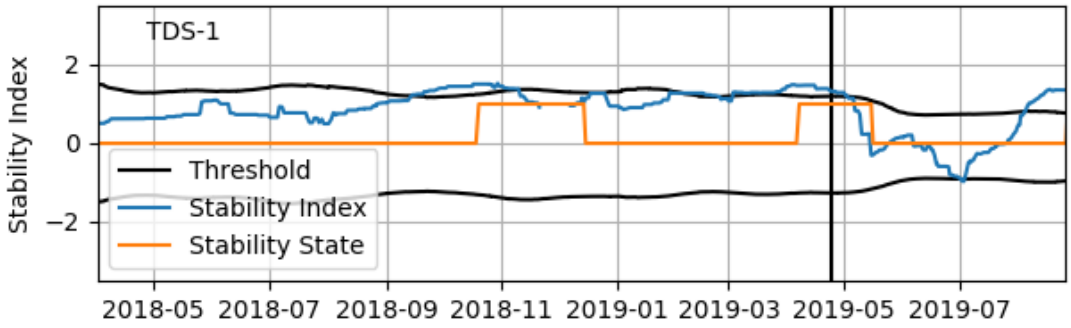

Fig. 7: Stability Index time series for the TDS-1 (NORAD ID 40076) satellite. The black curves are the 10% false-positive thresholds for this data set, and the orange line illustrates the attitude stability classification state (low indicating no evidence of stabilization, high indicating attitude stabilization). The DOS deploy event on April 24, 2019 is indicated by a solid, vertical black line.

4.2 Identification of Iridium Satellites

Iridium Communications has launched nearly 200 satellites for global telecommunications services. Their original constellation was launched between 1997 and 2002, and they launched a newer constellation (called Iridium NEXT) between 2017 and 2019. Of particular interest to this study, the two constellation are significantly different in physical shape.

Fig. 8 shows a scatter plot of the median RCS for each Iridium satellite vs. the mean SI for that satellite. The original constellation is shown in blue, while the NEXT constellation is shown in orange. There is a striking separation between the two constellations on the SI axis, with the NEXT constellation showing very strong stability signal.

We remind the reader that the lack of stability signal for the original constellation does not imply a lack of stability control for these satellites. An SI of zero can also indicate a stabilized satellite with near-isotropic RCS.

4.3 Attitude Stability of Orbiting Rocket Bodies

LeoLabs tracks more than 700 RSOs in low-Earth orbit classified by CSpoC as rocket bodies. These are typically vehicle upper stages that were left in orbit after delivering their payloads. Many have been in orbit for decades.

Rocket bodies are not expected to have active attitude control after the payload has been delivered. Therefore, one might suppose that they have tumbling motion. Fig. 9 shows our median RCS vs. mean SI for 738 rocket bodies. Although the vast majority of the SIs are close to zero, there is an intriguing tale to the distribution with positive SIs. This indicates that a fraction of rocket bodies may, in fact, be in a state of passive stabilization. This is consistent with theoretical studies of rocket body spin dynamics predicting that they may reach a state of passive stabilization[6][2].

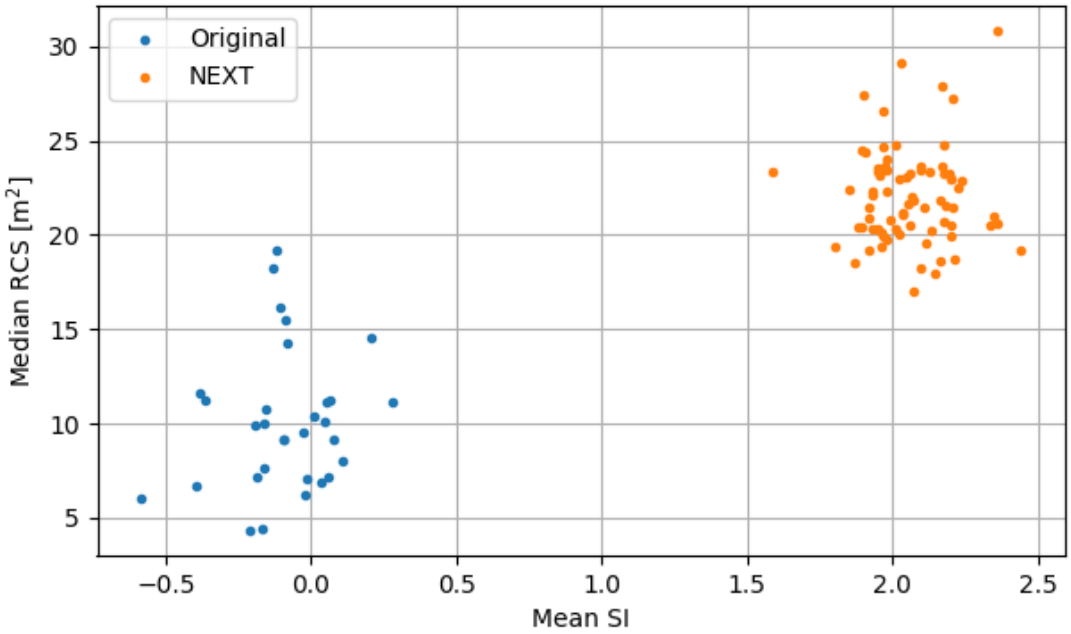

Fig. 8: Median RCS and mean SI values for the Iridium satellites. The original constellation is shown in blue, while the NEXT constellation is orange. There is a clear separation between the two constellations.

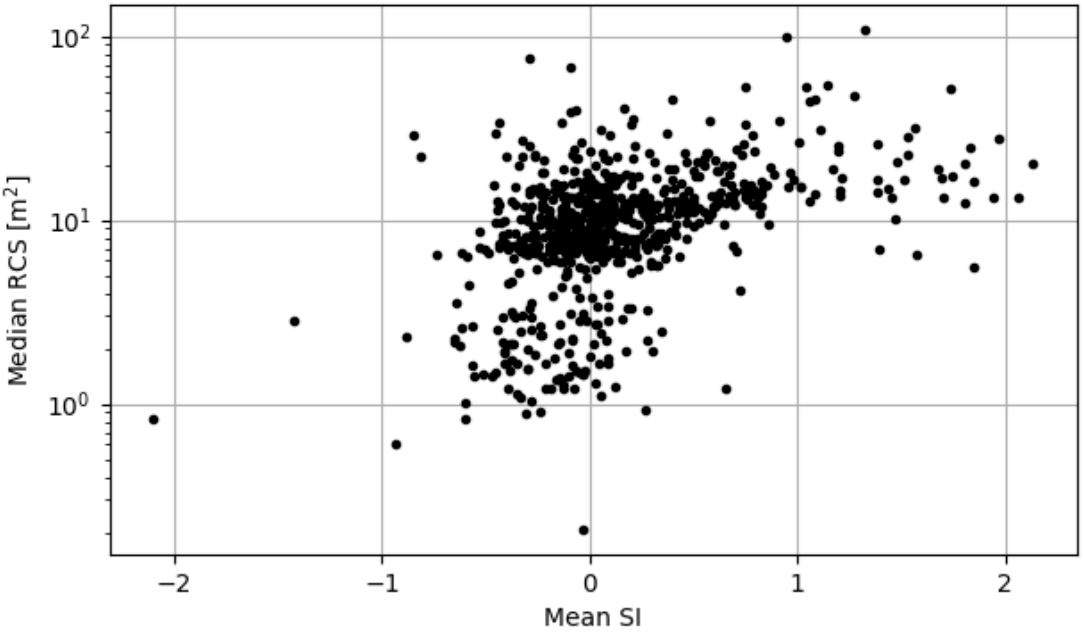

Fig. 9: Median RCS and mean SI values for rocket bodies in LEO. Most of the objects have SI values near zero, however there is a significant number with high SI values. This suggests that some rocket bodies have reached a state of passive stabilization.

4.4 Identification of Cosmos Constellations

Since 1962, the Soviet Union and Russian Federation have launched more than 2500 satellites with the designation Cosmos. These satellites have had a wide variety of purposes and orbits. Many of the satellites have since become inactive, but remain in orbit. One might naively expect the inactive satellites to eventually enter a state of tumbling motion. However, many of the Cosmos satellites used passive gravity-gradient stabilization. Assuming that the gravity-gradient booms of those satellites remain intact, it is likely that many of the older Cosmos satellites continue to be attitude stabilized.

Fig. 10 shows our median RCS vs. mean SI for the Cosmos satellites currently in LEO. In the top panel, the satellites have been color-coded to show to which constellation they belong. There are clear differences in the distributions for the varioius constellations. The bottom panel is color-coded for those constellations which are known to be gravity-gradient stabilized. It is notable that the gravity-gradient stabilized constellations tend to show higher SIs than those that are not gravity-gradient stabilized. There are some gravity-gradient stabilized constellations that do not show high SIs, and this is consistent with those satellites having isotropic RCS distributions. This demonstrates that the RCS statistics, including the SI, have clear value for characterizing satellites in constellations.

5. FUTURE PERFORMANCE IMPROVEMENTS

The results presented in this paper are using RCS data from a single radar site. We are working to augment this analysis through the addition of RCS data from other radar sites. We expect the performance to improve through two mechanisms: increased numbers of RCS measurements, and RCS measurements at different frequencies.

With additional radar sites, a given RSO will pass over our radar more times per day. This will provide a larger number of RCS measurements for each RSO. This will allow us to decrease the 60-day period for SI calculation and to decrease the SI thresholds, while maintaining a low false-positive rate. This will allow us to detect stability signatures from RSOs with less-pronounced RCS anisotropy or with frequent, controlled attitude changes.

LeoLabs' next generation of radar sites operates at higher frequencies than the current generation. RCS is frequency-dependent, so the RCS isotropy of a given RSO may be different at the two radar frequencies. This implies the possibility that the SI for different RSOs may be more or less sensitive to attitude stability when measured at different radar frequencies. We therefore expect to broaden our attitude stability sensitivity as we add more radar frequencies to our measurement suite.

6. CONCLUSION

We have demonstrated a new method for using RCS measurements to characterize the attitude stability of RSOs via a Stability Index. Through a careful process of data characterization and Monte Carlo simulations, we have successfully applied thresholds to this SI. This has allowed us to label RSOs in our data set with a stabilization classification.

In this paper, we have demonstrated the use of our Stabiltiy Index on a number of individual case studies. We are able to detect changes in attitude of active satellites, separate different constellations based on RCS and SI statistics, and detect passive stabilization of non-paylod RSOs and inactive satellites.

We expect this capability to continue to improve as LeoLabs continues to build our dedicated radar network.

7. REFERENCES

[1] European Space Agency MASTER-2009 Model. `https://sdup.esoc.esa.int/master/`.

[2] S Efimov, D Pritykin, and V Sidorenko. Long-term attitude dynamics of space debris in sun-synchronous orbits: Cassini cycles and chaotic stabilization. *Celestial Mechanics and Dynamical Astronomy*, 130(10):25, 2019.

[3] C Früh, M Jah, E Valdez, P Kervin, and T Kelecy. Taxonomy and classification scheme for artificial space objects. *Proceedings of AMOS*, 2013.

[4] D. Hall, B Calef, K Knox, M Bolden, and P Kervin. Separating attitude and shape effects for non-resolved objects. *Proceedings of AMOS*, 2007.

[5] MJ Holzinger, KT Alfriend, CJ Wetterer, KK Luu, C Sabol, K Hamada, and W Harms. Attitude estimation for unresolved agile space objects with shape model uncertainty. *Proceedings of AMOS*, 2012.

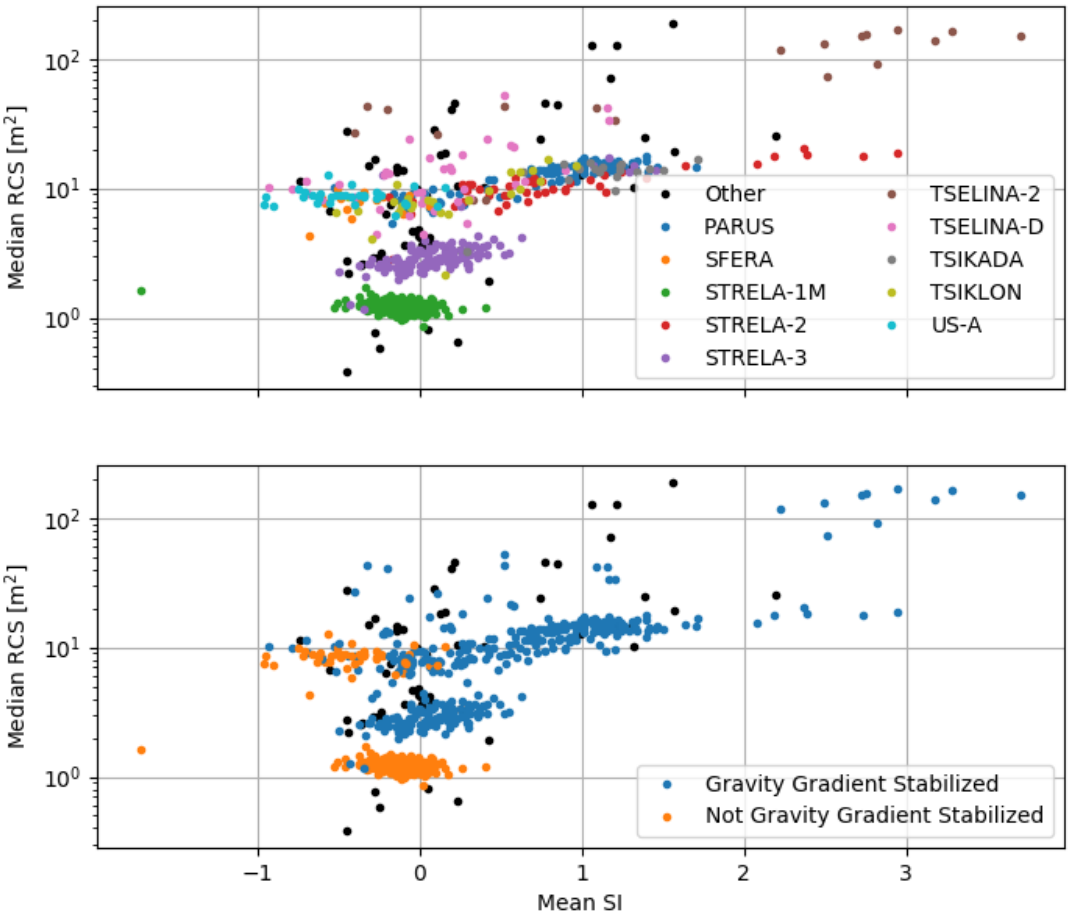

Fig. 10: Median RCS and mean SI values for Cosmos payloads in LEO. In the top panel, the objects have been color-coded by parent constellation. The bottom panel shows the same objects, but color-coded by whether or not their parent constellation had a gravity-gradient stabilization design. The top panel illustrates that the various constellations have distinct RCS vs. SI properties, while the bottom panel illustrates the trend of higher SI values for gravity-gradient stabilized satellites.

[6] GW Ojakangas, P Anz-Meador, and H Cowardin. Probable rotation states of rocket bodies in low earth orbit. *Proceedings of AMOS*, 2012.

[7] MA Richards. Target fluctuation models. In MA Richards, JA Scheer, and WA Holm, editors, *Principles of Modern Radar*, chapter 7, pages 247–271. SciTech Publishing, Raleigh, NC, 2010.

[8] LBM Sàgnieres and I Sharf. Long-term rotational motion analysis and comparison to observations of the inoperative envisat. *Journal of Guidance, Control, and Dynamics*, 42(2):364–376, 2019.