Monitoring Satellite Pattern-of-Life Changes with Passive Radio Frequency Data

Harris Mohamed
Kratos Defense

ABSTRACT

As space evolves into a critical infrastructure sector, it is becoming more congested and more contested. As a result, Space Domain Awareness (SDA) is significantly elevating in importance. In addition to Electro-Optical (EO), Radio Frequency (RF), and Ground-based radar positioning information, another critical SDA component is satellite payload behavior. Given these developments, this paper focuses on establishing the pattern-of-life for space objects using RF data.

While EO and radar have their respective strengths, they each also have critical shortcomings. For instance, optical observations can only be made at certain times of the day and in ideal weather conditions. Similarly, radar takes a vast amount of power to track objects in Geosynchronous Orbit (GEO). Furthermore, radar presents environmental and regulatory concerns given its active nature. Both approaches primarily focus on the orbital state of the object, while neither provide much insight into payload behavior. This leaves a significant information gap for the SDA observers using these systems.

These optical and radar-based system limitations can be overcome for objects that transmit RF energy because RF data can be collected in all weather conditions and at any time of day. The RF measurements collected from these objects can enhance the SDA mission by providing additional insight into the behavior and operations of both the space objects' payloads as well as any signals they may be transmitting. Additionally, RF characterizations of downlinked signals (e.g., power, frequency, modulation type, data rate, etc.) can be taken at regular intervals for all observed signals. This provides SDA observers with a deeper understanding of the payload, to include determining intent, attribution, and detecting deviations in behavior.

To accurately perform pattern-of-life analysis on RF data, this paper will explore a variety of algorithms that broadly fall into three categories: statistical analysis, time series analysis, and unsupervised machine learning approaches. RF measurements collected from an object can be represented as both univariate and multivariate time series. First, classical statistics methods will be used to establish baseline performance by flagging anomalous measurements that fall outside 3 standard deviations. Then, we will apply state-of-the-art time series analysis algorithms—such as Autoregressive Integrated Moving Average (ARIMA) and Ruptures[1] (a change point detection algorithm)—to the RF time series. Finally, unsupervised machine learning algorithms, such as LSTM Auto-Encoder (LSTM-AE), will be applied to perform anomaly detection. The best of these approaches will be combined in an ensemble method used to collectively vote on anomalies. These approaches will be tested on real-world data and used to prototype a system that can autonomously alert operators of suspicious behavior on an object's payload.

In this paper, we convey the value that RF pattern-of-life measurements can provide to the SDA mission. First, we introduce the basic concept of operations for collecting and using RF measurements of a space object and its payload in addition to describing the various measurements that can be made. Then, we demonstrate how these basic measurements can be accumulated over time to build a pattern-of-life estimate for space objects as well as discuss approaches for using changes in this pattern-of-life to identify events relevant to SDA missions. Lastly, we evaluate the accuracy of statistical, time series, and machine learning algorithms to combine them into a pattern-of-life ensemble model. The pattern-of-life ensemble model will be applied to several examples. Nominal real-world data will be used to demonstrate the capabilities of the pattern-of-life ensemble model. Finally, the ensemble model will be applied to anonymous, real-world data that was captured during the jamming of commercial satellites used by Russia during the Russia/Ukraine conflict.

1. INTRODUCTION

Traditional approaches to behavioral Space Domain Awareness (SDA) have primarily implemented Electro-Optical (EO) and radar measurements. Light curve analysis has used EO to uniquely identify satellites and to detect satellites tumbling. Likewise, radar cross-section analysis is used to determine satellite size. While both phenomenologies have their respective strengths, each also has critical shortcomings.

For instance, optical observations can only be made at certain times of the day and in ideal weather conditions. Similarly, radar takes a vast amount of power to track objects in Geosynchronous Orbit (GEO). Furthermore, radar presents environmental and regulatory concerns given its active nature. Both approaches primarily focus on the orbital state of the object, while neither provide much insight into payload behavior. This leaves a significant information gap for the SDA observers using these systems.

These optical and radar-based system limitations can be overcome for objects that transmit Radio Frequency (RF) energy. This is because RF data can be collected in all weather conditions and at any time of day. The RF measurements collected from these objects can enhance the SDA mission by providing additional insight into the behavior and operations of the space objects' payloads as well as any signals they may be transmitting. Additionally, RF characterizations of downlinked signals (e.g., power, frequency, modulation type, data rate, etc.) can be taken at regular intervals for all observed signals. This provides SDA observers with a deeper understanding of the payload, such as determining intent, attribution, and detecting deviations in behavior.

While point observations of RF measurements can be useful in understanding the current state and behavior of an object, these measurements can also be accumulated over time to develop a pattern-of-life for both the space object and its payload. RF pattern-of-life observations enable new levels of understanding that were not previously available through radar or optical technologies. Once a pattern-of-life is established, SDA systems can identify deviations from normal states as indicators of behavior that may require additional attention or observations. For example, changes in an object's assigned bandwidth usage could serve as an indicator for changes in real-world events (i.e., political, economic, or military). Furthermore, any large and unexpected RF metrics that deviate from the pattern-of-life observations could be potentially attributed to malicious intent, such as jamming.

One of the difficulties with pattern-of-life analysis on RF data is that the data is unlabeled. Traditional anomaly detection workflows will operate on labelled data, which is data that has tags to help identify it as anomalous or nominal. Once an algorithm is trained to correctly classify anomalous data, it can be run on new data not yet seen by the model. Literature typically refers to this as supervised training. RF datasets are extremely large, which makes it difficult for any person to label datasets. Furthermore, what is considered anomalous for one satellite may not be anomalous for another. To add to the complexity, every RF analyst will have different expectations from the anomaly detection algorithms. For instance, an analyst might want an anomaly detection algorithm to be very sensitive to changes on a Russian payload but may not desire the same thing out of an American communications satellite. The challenge is to design a system which is functional, scalable, and customizable.

The following sections provide an overview of RF concepts and explain the various techniques (i.e., statistical, time series, and machine learning) that can be applied to this data. The final section of this paper describes the ensemble model and then applies it to a few examples.

2. RADIO FREQUENCY OVERVIEW

When operating in the RF domain, there are several terms and concepts that are important to understand.

- *Transponder*: A satellite communication unit that receives data from a ground antenna and then transmits that data to a different ground antenna. Satellites will typically have dozens of transponders, with each on a different portion of the RF spectrum. Transponders are defined by their frequency range, beam connectivity, and polarization.

- *Carrier*: A signal that a transponder emits. Carriers are modulated signals that carry data between ground stations. A transponder supports one or more carriers.

- *Payload Usage*: For every RF capture of a transponder with bandwidth b_t, there are n detected carriers, each with a bandwidth b_c.

  The Payload Usage can then be defined as $\frac{\sum_{i=0}^{n} b_{c,i}}{b_t}$

  (Note that this formula does not take guard band into account.)

  Fig. 1 displays a transponder wherein all traffic stops and causes the payload usage to drop to 0.

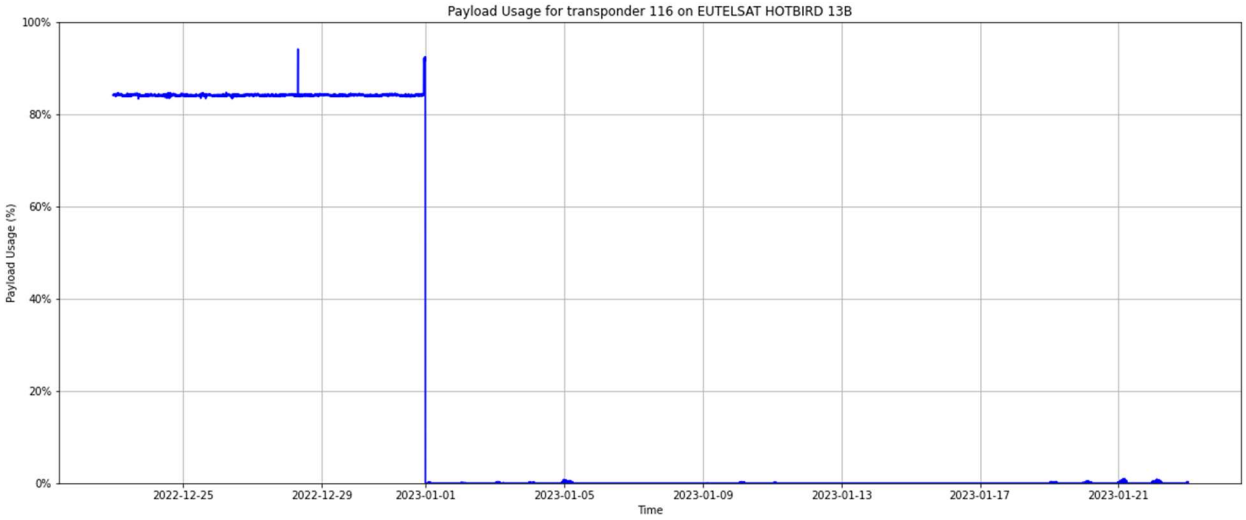

Fig. 1. Payload Usage Plot for Transponder 116 on Eutelsat Hotbird 13B wherein All Traffic Stops on 2023-01-01

- *Waterfall*/*Spectrogram*: For a transponder, this visualizes the RF spectrum using three dimensions: Frequency (typically in MHz), Time, and Power (typically in dBW). The terms "waterfall" and "spectrogram" are generally interchangeable terms. In this work, we will refer to these diagrams as waterfalls for the sake of consistency. Waterfalls are like heatmaps that give an analyst an at-a-glance view of payload behavior over time for an entire transponder. Fig. 2 shows a waterfall that has 7 distinct carriers, with each carrier designated by a horizontal bar. (Note that for clarity, the x-axis is time, and the y-axis is frequency.)

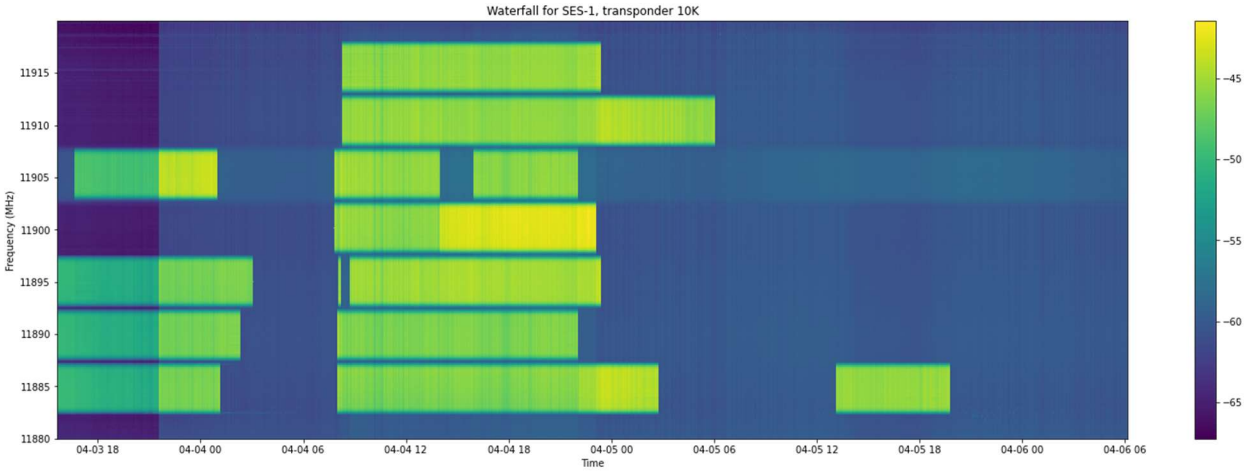

Fig. 2. Waterfall Plot of Transponder 10K on SES-1

- *Power Spectral Density (PSD):* Another way to visualize RF data. The x-axis is the RF frequency, and the y-axis is the signal power (typically in dBW). This can give an RF analyst an at-a-glance view of instantaneous payload behavior for a transponder. Fig. 3 displays the PSD for an SES-1 transponder with 7 carriers at a particular instance of time.

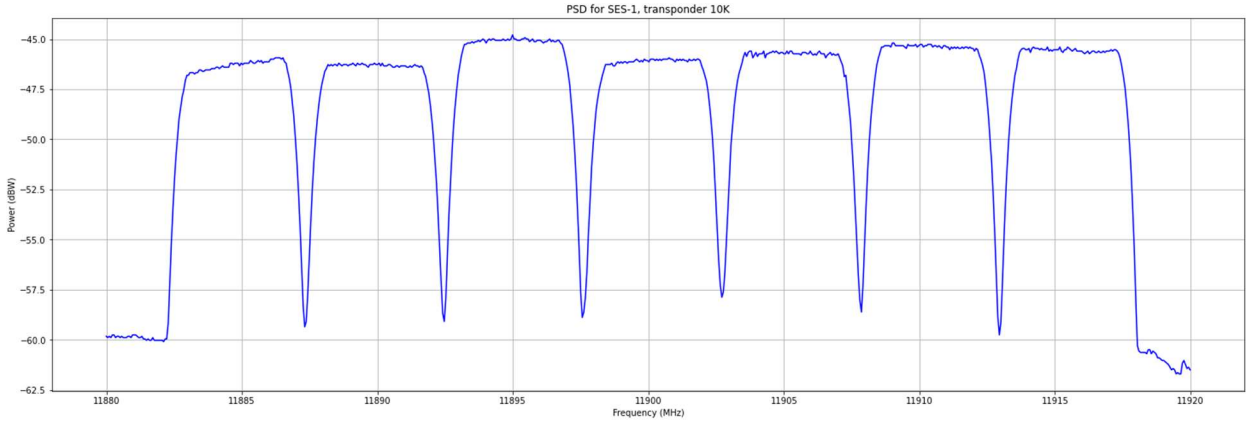

Fig. 3. PSD for SES-1, Transponder 10K on SES-1 with 7 Carriers

When transponder and carrier information is known, it can be visualized as shown in Fig. 4.

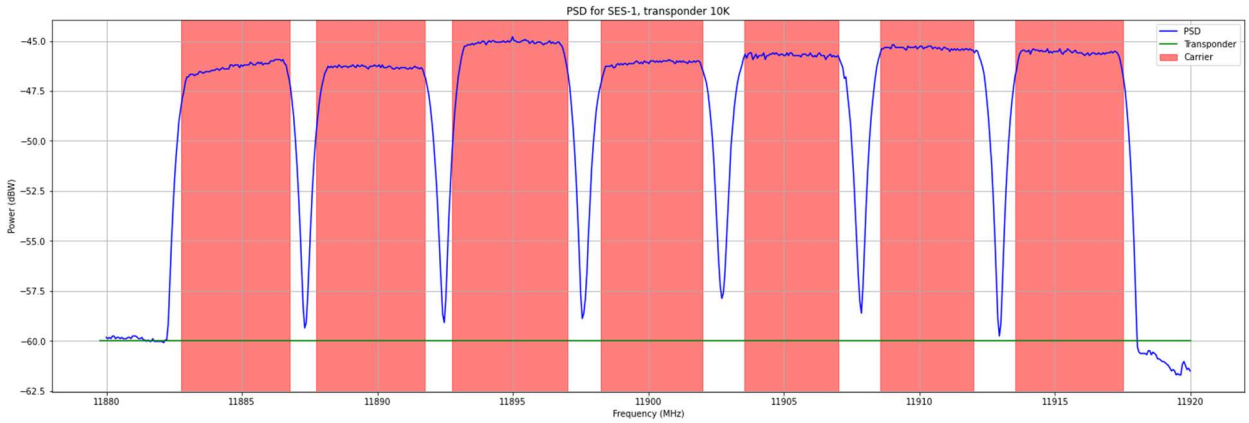

Fig. 4. PSD of a Transponder with Transponder (green) and Carriers (pink) Highlighted

- *Symbol*: For encoded carriers, the modulation process encodes bits of data into distinct RF states called symbols. The modulation type determines how many bits of information are contained in each symbol. For example, a carrier with modulation type 8PSK will encode every 3 bits into a single RF symbol. This knowledge is used by DSP algorithms to convert a stream of symbols back to the original transmitted message.
- *Symbol Rate*: The transmission rate of signals on a connection in symbols per second. This is usually represented in kilosymbols per second (ksps) or megasymbols per second (msps).
- *Carrier-to-Noise (C/N) Ratio:* Higher ratios indicate better reception quality and in general, higher communications accuracy and reliability.

For a carrier, this is defined as $\frac{Power_{Carrier}}{Power_{Noise}}$

Energy per Symbol-to-Noise Density (Es/N0) Ratio: Es/N0 is similar to C/N, with the exception that it measures the normalized ratio of the power of the symbol to noise. This metric is of particular importance because it is determined on the receive side after the signal is demodulated.

- *Interference*: There are many different kinds of carrier interference, but this work formally defines it as

$$\frac{C}{N} - \frac{Es}{N0}$$

Nominally, this should be ≤ 1 for carriers. However, when a carrier has interference, the C/N will increase and the Es/N0 will decrease. This derived metric is a good candidate for anomaly detection because it can be used to find potential interference autonomously.

3. PATTERN-OF-LIFE ANALYSIS TECHNIQUES

A pattern-of-life can be established once enough data is collected from a space object transmitting RF. The longer the time between RF collections, the worse the pattern-of-life analysis becomes. All the algorithms described in this work are anomaly detection algorithms. The deviations from a pattern-of-life are what an RF analyst will be alerted on. To this end, this paper evaluates three pattern-of-life analysis technique categories: Statistical, Time Series, and Machine Learning.

4. STATISTICAL PATTERN-OF-LIFE ANALYSIS TECHNIQUES

This section describes the application of classical statistical algorithms to the RF time series for anomaly detection. Although limited, these statistical analysis methods establish the baseline performance for the more sophisticated techniques detailed in the following two sections.

- *Sliding Gaussian (Naïve Method):* This is the simplest approach because it compares the current data point (c) to the previous data point (p) with a specified time lag (l). If the difference is greater than a specified number (n) of standard deviations (σ), then c is marked as an anomaly.

  Formally, the condition for anomaly detection can be expressed as $(c-1) \geq n*\sigma$

  This performs quite well on clean datasets because it clearly marks the abrupt changes in payload usage. However, on noisier datasets, this approach is too sensitive to tiny changes. For example, Fig. 5 shows this approach applied to a clean dataset, and Fig. 6 shows it applied to a noisy one (note that boundaries are removed from this plot for clarity).

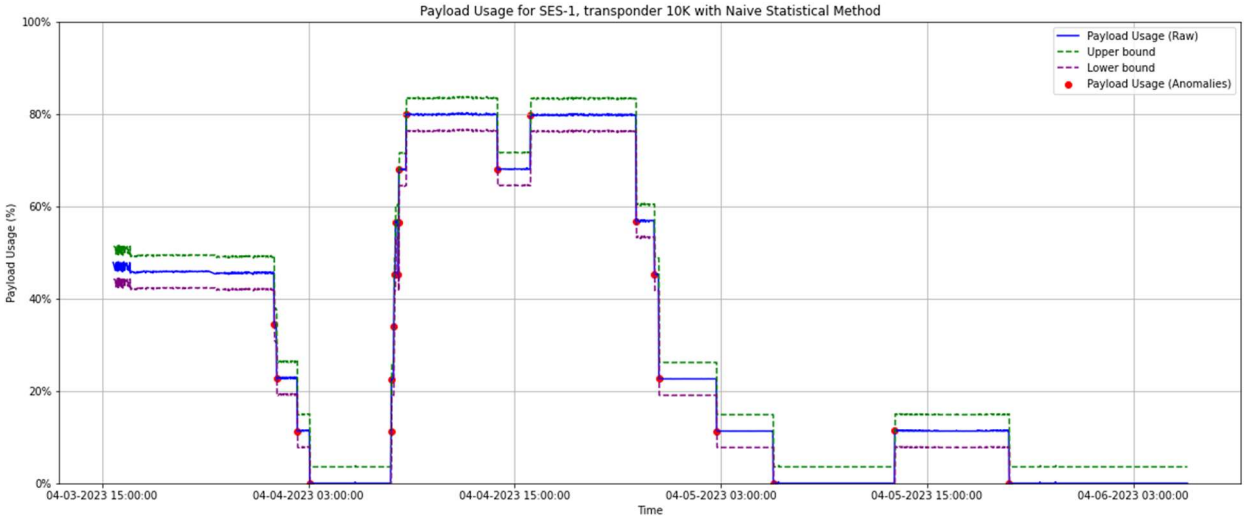

Fig. 5. Naïve Method Run on Clean Data (l = 1, n = 3)

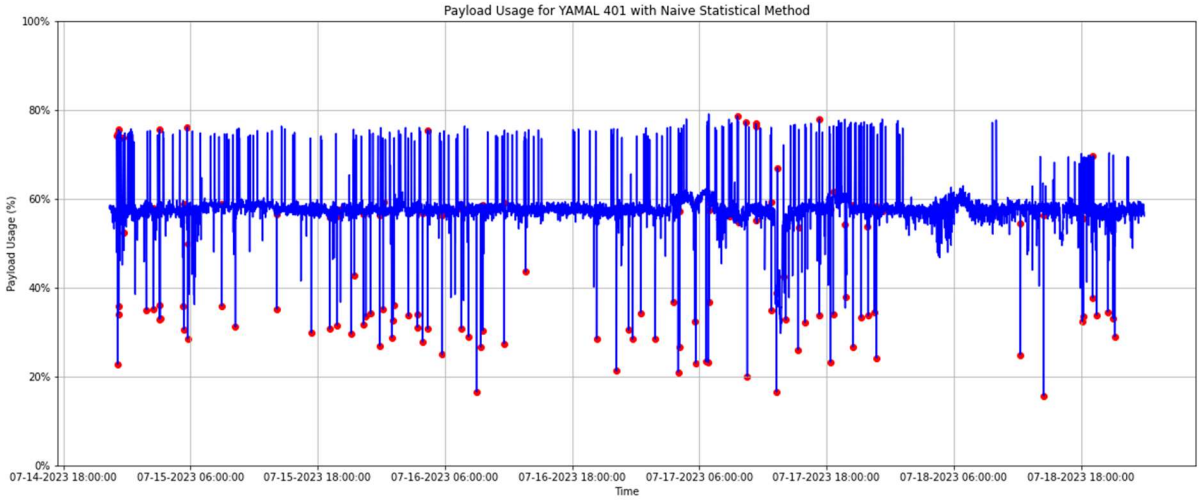

Fig. 6. Naïve Approach Run on a Noisy Dataset (l = 1, n = 3)

- *Sliding Gaussian + Moving Average (MA)*: This is similar to the Sliding Gaussian (Naïve Method) approach. However, rather than processing the raw data, we smooth the data using an MA with a window size s. The Sliding Gaussian is then applied with a window size w. Fig. 7 displays how this performs better than the Sliding Gaussian (Naïve Method). Because most of the noise is smoothed out, we only detect anomalies for the most abrupt changes in the payload usage rather than all the local extrema from the noisy data.

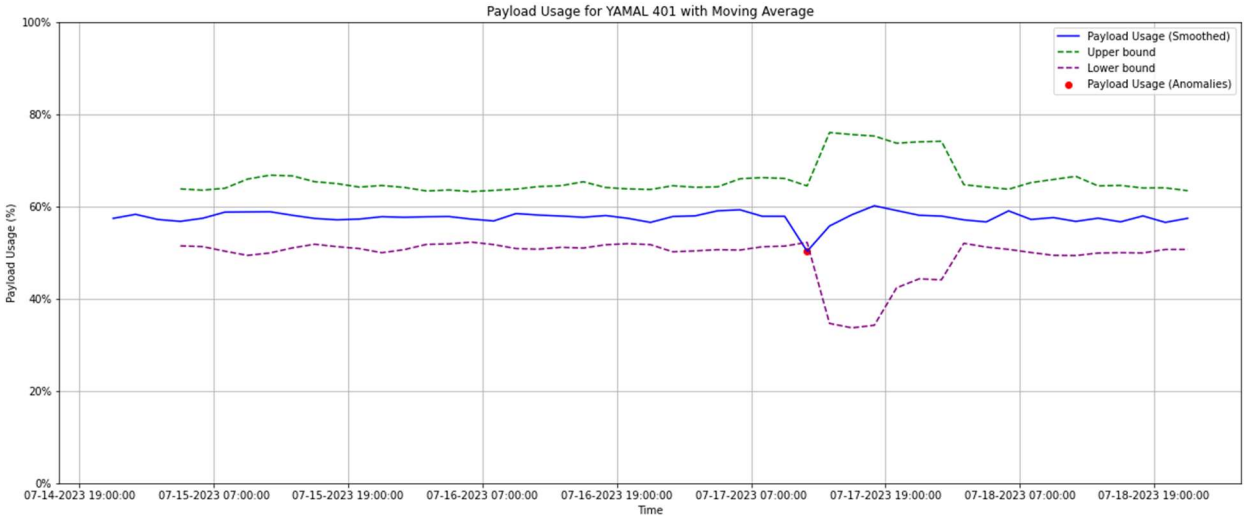

Fig. 7. Sliding Gaussian Applied to Payload Usage for YAMAL 401 Smoothed with MA (s = 2 hours, w = 6 hours, n = 3)

- *Sliding Gaussian + Moving Median*: This is similar to the Sliding Gaussian + MA approach, with the addition of using a Moving Median with a window of specific size (w) as the smoothing function. The Moving Median is commonly used in statistics to remove outliers. For this dataset, however, it smooths the data too much. Therefore, the overall trends of the dataset are masked by the median. Fig. 8 shows an example of no anomalies detected with the same parameters.

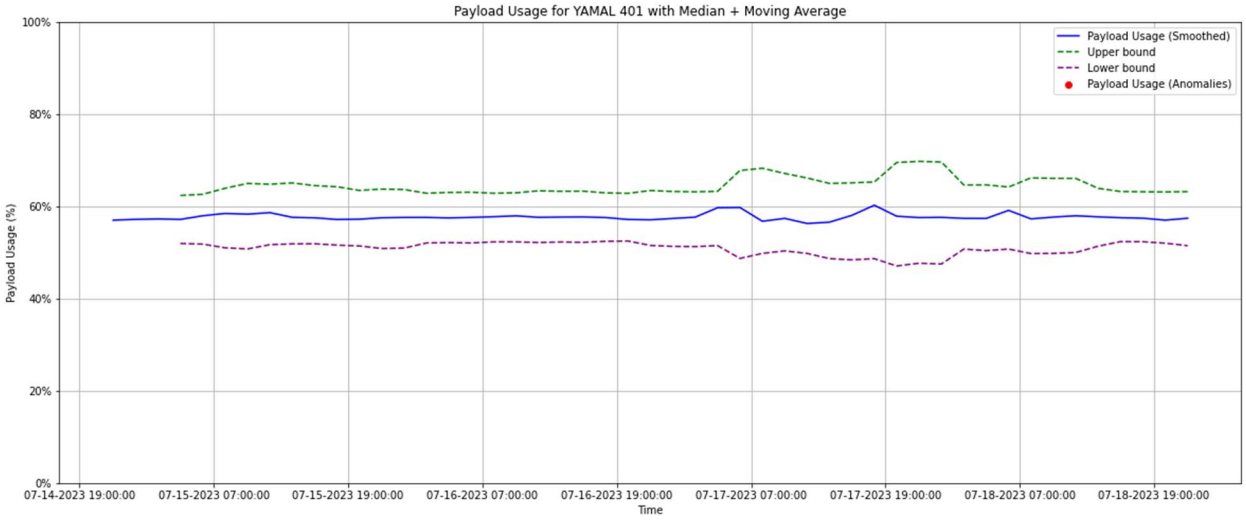

Fig. 8. Sliding Gaussian Applied to Payload Usage for YAMAL 401 Smoothed with Moving Median (s = 2 hours, w = 6 hours, n = 3)

- *Sliding Gaussian + Moving Exponential Average:* In this approach, we first apply a moving exponential weight filter to the raw data. This gives more weight to recent observations and less weight to past observations according to a smoothing factor, α. We then apply a moving average with a sampling window s to the data. Together, this smooths short-term fluctuations and enables us to observe long-term trends. Compared to the previous statistical approaches, this Sliding Gaussian + Moving Exponential Average approach smooths the data enough to not be overly noisy, while still retaining the behavior of the raw dataset reasonably well. Fig. 9 shows an example dataset with the Moving Exponential Average applied to the payload usage data.

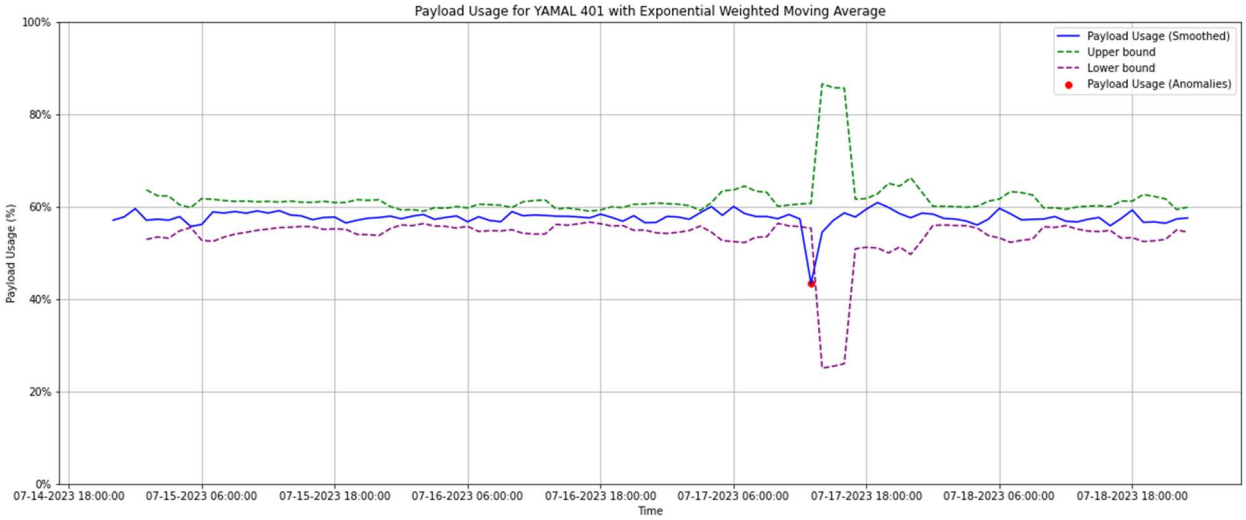

Fig. 9. Moving Exponential Average Applied to Payload Usage Data for Yamal 401 (α = 0.15, s = 2 hours, w = 6 hours, n = 3)

- *Sliding Interquartile Range (IQR):* This approach calculates the IQR for a selected window size of data w, and then identifies which points in the future lie outside the IQR. The data is smoothed using the Moving Exponential Average with a window size s. IQR tends to be too sensitive to small changes in the time series. Fig. 10 shows an example dataset with the Sliding IQR applied to the payload usage data.

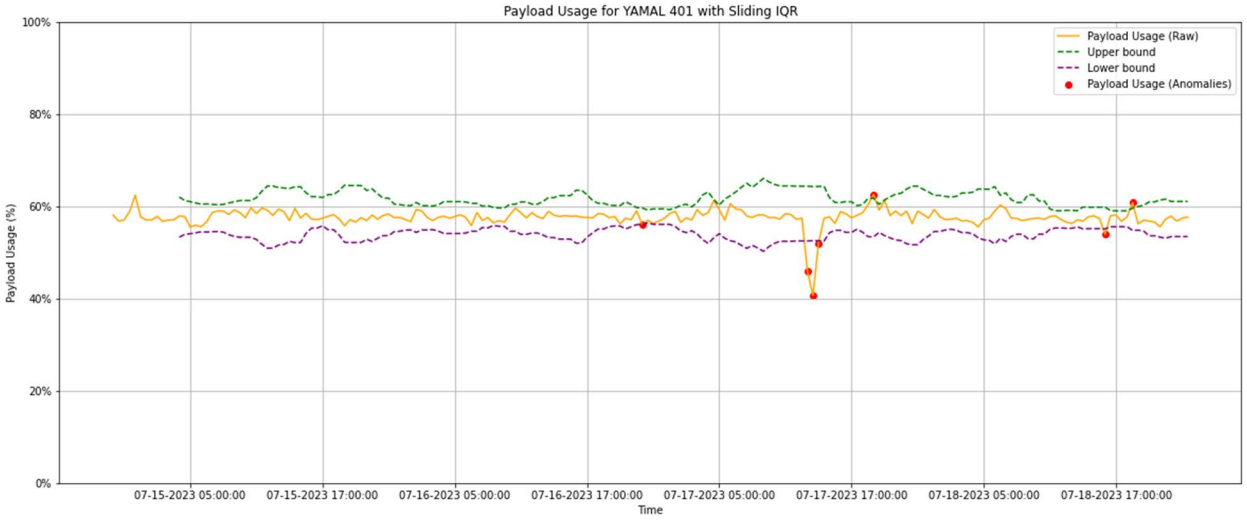

Fig. 10. Sliding IQR Applied to Payload Usage Data for Yamal 401
(s = 2 hours, w = 6 hours, n = 3)

- *Sliding Minmax:* This approach smooths data using the Moving Exponential Average. The maximum and minimum values for windows of size w are then calculated across the dataset. Next, the extrema are multiplied by a certain amount δ to create the upper and lower bounds for the data. The calculated bounds are then applied to a certain number of points in the future, and the points that lie outside of the boundaries are flagged as anomalies. Fig. 11 displays how well the Sliding Minmax approach captures the anomalies.

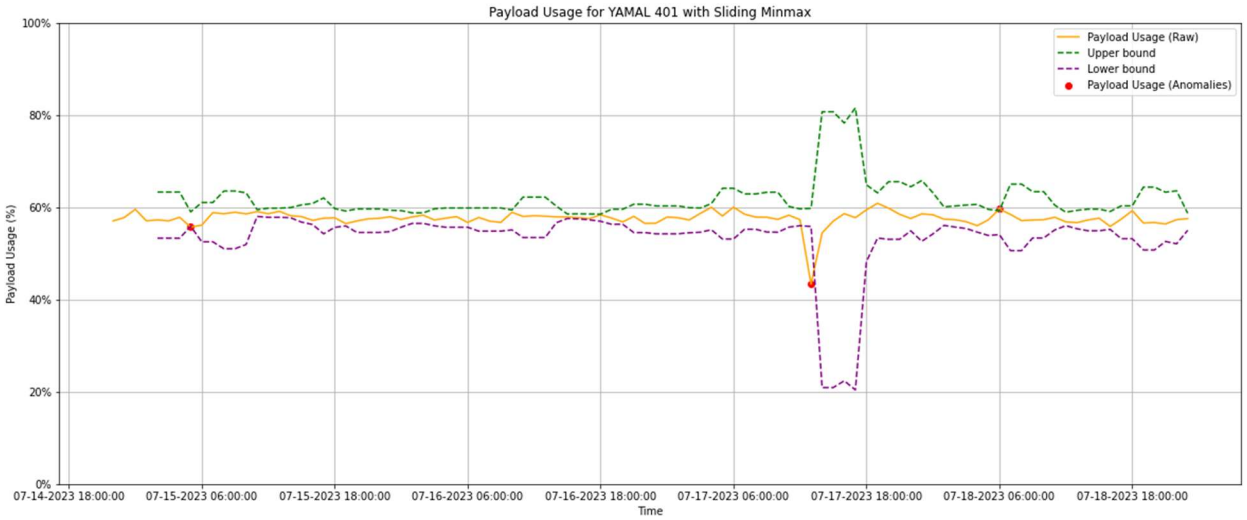

Fig. 11. Sliding Minmax Applied to Payload Usage Data for Yamal 401
(w = 1 hour, δ = 0.01)

Each analysis technique described herein has its own strengths and weaknesses. For example, the Moving Exponential Average is the most general anomaly detection method of all the statistical techniques. Due to its ease of customization and simplicity, it is the best approach for an RF analyst to apply at scale. The Sliding IQR approach is also good given that it runs fast and is easy to tune parameters according to the use case. Ultimately, it is up to the RF analyst to discern which approach is most appropriate to use depending on the specific scenario.

5. TIME SERIES PATTERN-OF-LIFE ANALYSIS TECHNIQUES

This section explores the application of well-established, state-of-the-art, time series pattern-of-life analysis methods for anomaly detection.

- *Autoregressive Integrated Moving Average (ARIMA) Model*: This is a time series technique that aims to forecast future values based on previous values. ARIMA Models use a wide variety of sophisticated statistics techniques to describe trends in a time series dataset. (Note that the exact techniques are beyond the scope of this paper [2]).

  For this technique, we feed the first 60% of the time series dataset and then forecast the remaining 40%. As shown in Fig. 12, the ARIMA model predicts the next data point and then retrains based on the error. If the forecasted value is drastically different than the actual value, then the actual value would be marked as an anomaly.

  This workflow could work well at scale in operation; however, in order for this approach to be feasible, ARIMA must be good at forecasting for nominal datasets. Furthermore, successfully training it involves analyzing a large volume of convoluted plots. There are software libraries that will automatically tune ARIMA, but the search takes too long to be practical. For these reasons, this analysis approach is unsuitable for RF analysts due to the steep learning curve and the time involved.

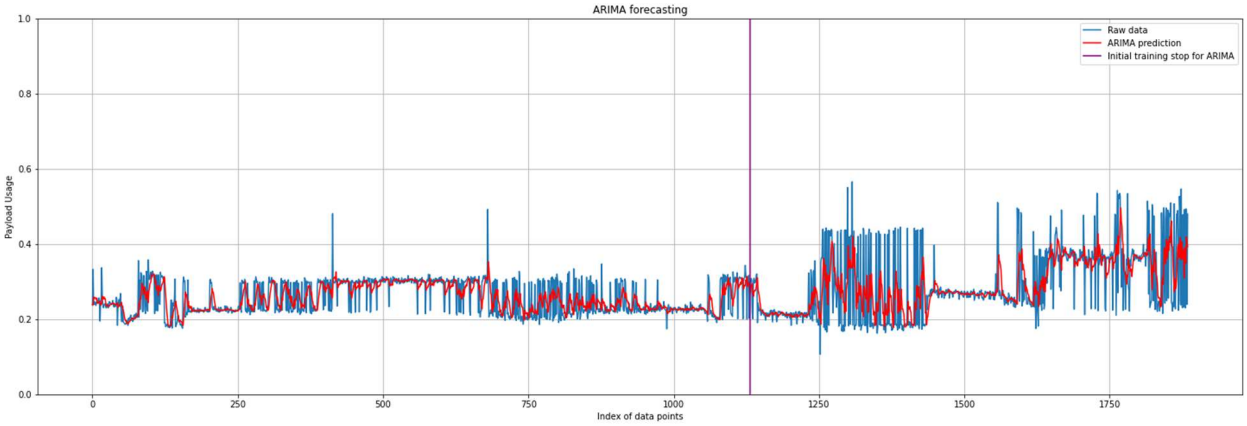

Fig. 12. ARIMA Forecasting for CHINASAT 11

- *Ruptures*: This is a state-of-the-art segmentation algorithm that aims to find change points in time series datasets. Rather than forecasting future points, this analysis approach takes the entire dataset into account and segments it as a means to identify points in time where the time series drastically changes. There are two methods of segmentation: Binary Segmentation and Pruned Exact Linear Time (Pelt) [1]. Within these parameters, there are only the penalty (p, threshold for change point detection) and jump (j, the grid of possible change points) values to tune.

  Both of these methods are fast: Binary segmentation runs in $n \log n$ time; and, with appropriately chosen hyperparameters, Pelt can run in linear time [3]. Due to the fast execution time and ease of tuning the parameters, these algorithms are suited to real-time application. Fig. 13 shows an example of the Ruptures analysis technique applied to payload usage data.

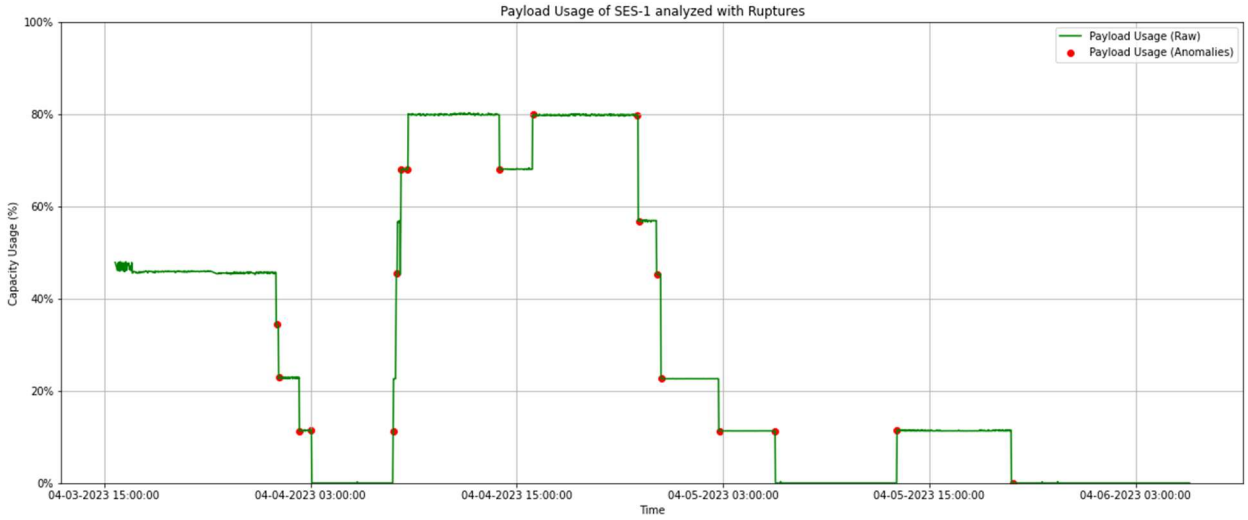

Fig. 13. Payload Usage Data from SES-1 Analyzed with Ruptures, Pelt ($p = 10, j = 5$)

Of these evaluated time series pattern-of-life techniques, we have determined that Ruptures is the best to implement due to its speed and ease of application. While ARIMA can achieve more nuanced anomaly detection, the number of parameters that must be tuned is quite high. Therefore, the slow run-time of automated ARIMA solutions means that it does not scale well to large datasets.

6. MACHINE LEARNING PATTERN-OF-LIFE ANALYSIS TECHNIQUES

This section explores two machine learning approaches for pattern-of-life analysis techniques.

- *Long Short-Term Memory (LSTM):* This technique is a very common neural network used on sequential data [4]. LSTM is particularly adept at learning patterns because it takes time into account during training. In the context of this paper, LSTM is trained on the first-third of the dataset and is then used to predict the remaining two-thirds. The means of predicting the second half of data is similar to that used for ARIMA, in which points are fed to the LSTM and then predicted forward. Within this pattern-of-life analysis, any predictions that deviate from the fixed threshold are classified as anomalies. Fig. 14 shows an example of the LSTM analysis techniques applied to payload usage data.

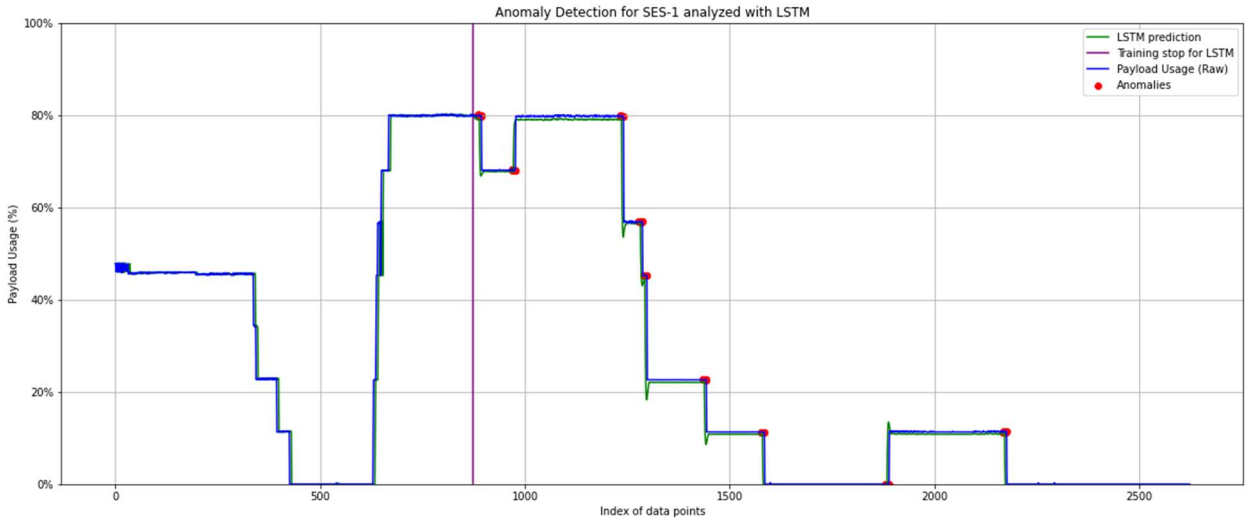

Fig. 14. LSTM Applied to SES-1 Payload Usage Data

- *Autoencoders + LSTM*: Autoencoders are a type of neural network with the capability to learn how to convert a high dimensionality dataset into a lower dimensionality representation. It learns efficient ways to code unlabeled data, which is suitable to this use case. An autoencoder consists of two functions: (1) an encoder that converts the higher dimensional data to a lower dimensional latent space, and (2) a decoder that converts the lower dimensional data from the latent space back to the input. These functions enable the input data dimensionality to be reduced in order to focus training on only the dimensions that matter.

  Using an autoencoder on its own would not work well for this dataset because of the time dependence involved. Traditionally, an autoencoder uses a dense, linear neural network to convert to and from the latent space. However, here we replace the dense neural network with an LSTM so the model can learn the time dependence. Any raw data points that deviate from a fixed threshold within the reconstructed time series are flagged as anomalies. Fig. 15 shows an example of the Autoencoder + LSTM analysis technique applied to payload usage data.

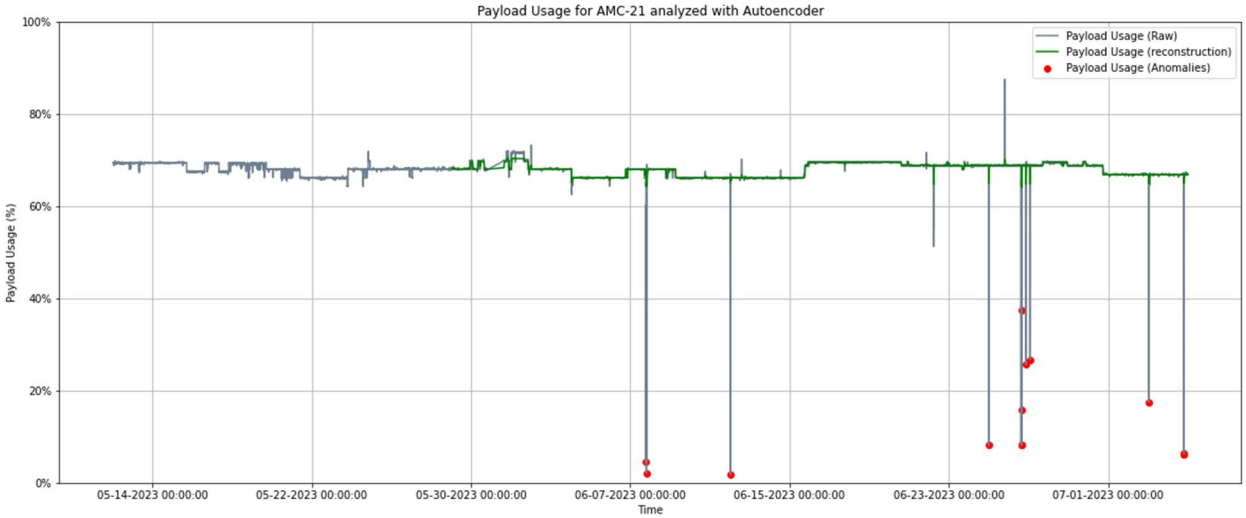

Fig. 15. Autoencoder + LSTM Applied to AMC-21 Payload Usage Data

7. ENSEMBLE METHOD

Now that several pattern-of-life analysis techniques have been defined, we will apply them to some scenarios. In previous sections, the exploration of statistical, time series, and machine learning approaches were visualized on transponder data. In this section, we will analyze carrier data by applying an ensemble method. This method applies the Sliding Gaussian + Moving Exponential Average, Ruptures, and Autoencoder + LSTM analysis techniques to the same dataset. Within this ensemble method, if two out of three methods identify and agree on an anomaly, it will be flagged as such.

Our first example demonstrates the ensemble model on is the Eutelsat Hotbird 13B, a communications satellite that provides direct-to-home broadcasting services. We analyze the payload usage on transponder 116. This example highlights one of the simplest kinds of anomalies that an RF analyst can identify—a complete drop in payload usage. While it is dangerous to speculate the reason why RF metrics change the way they do, we can use all the metrics available to derive an informed conclusion.

Fig. 16 shows the ensemble model detecting the drop in payload usage as an anomaly, as expected. Since the payload usage drops on January 1st, this could be an indication of a service contract ending. However, disparate data sources might offer alternative conclusions. For example, it could also be transponder failure, business ceasing, or a transponder maintenance window.

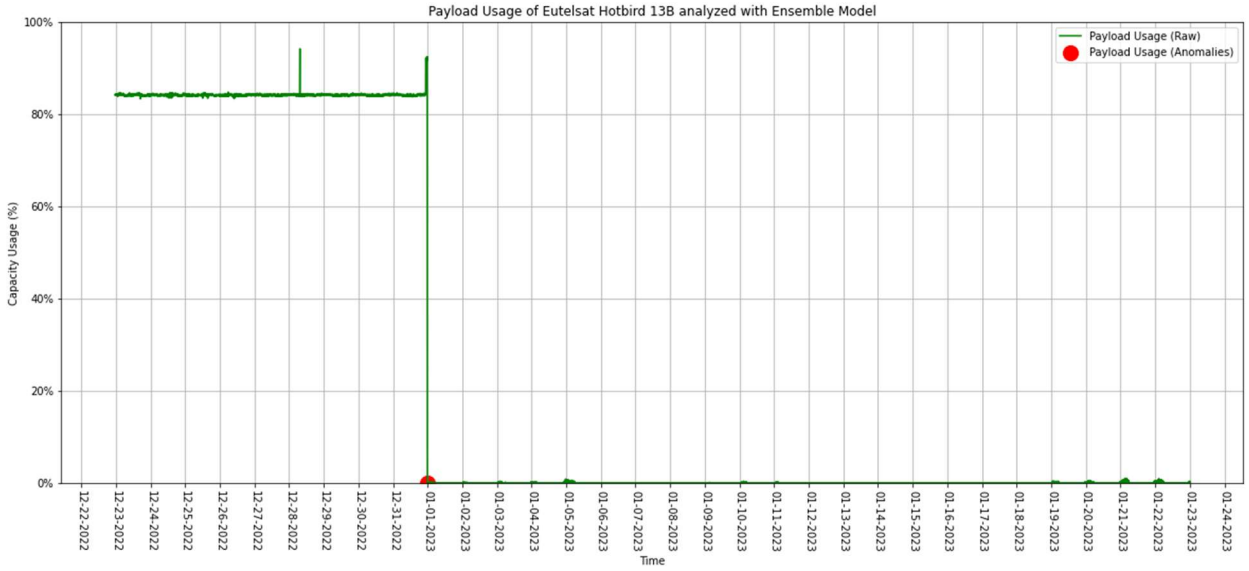

Fig. 16. Ensemble Model Applied to Eutelsat Hotbird 13B, Transponder 116

Fig. 17 displays that this transponder was comprised of a single carrier that disappeared on January 1st, 2023.

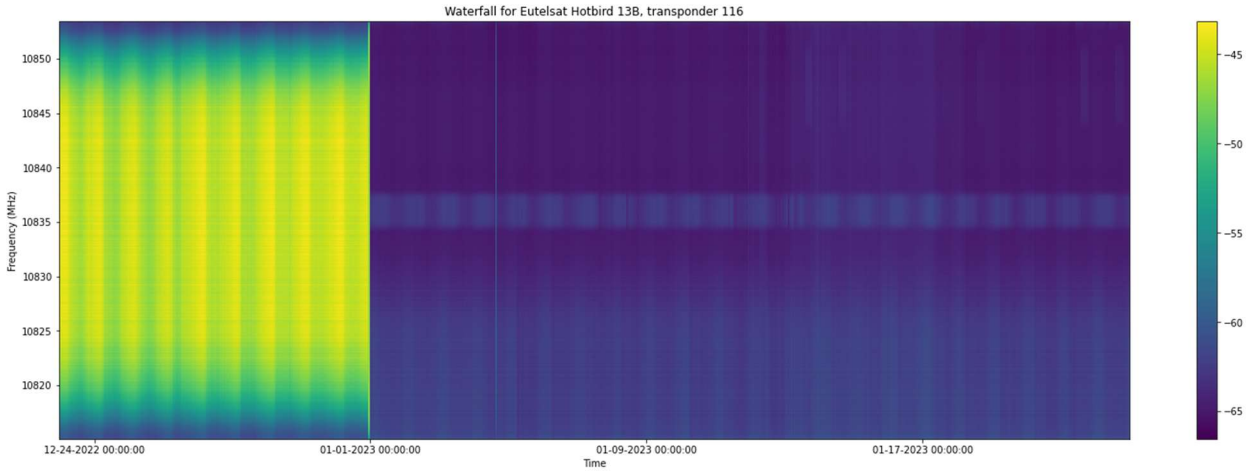

Fig. 17. Waterfall for Eutelsat Hotbird 13B

The next example is from SES-1, a communications satellite. As shown in Fig. 18, we applied the ensemble model on transponder 10K. In just 3 days, there were a large number of anomalies in the payload usage, indicating that something odd might be happening on this transponder.

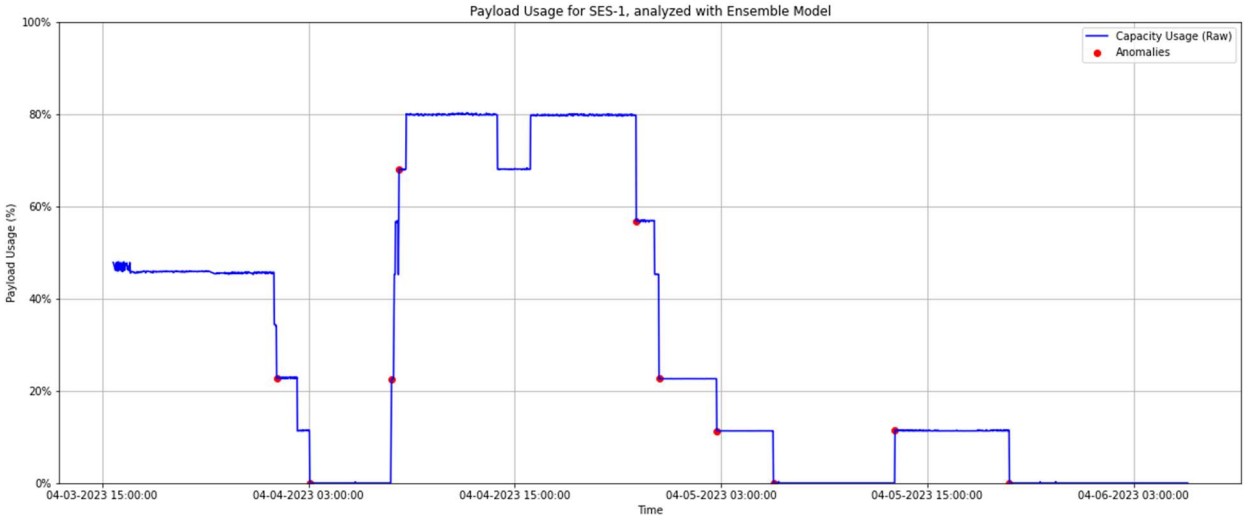

Fig. 18. Ensemble Model Applied to SES-1, Transponder 10K

The waterfall in Fig. 19 shows several carriers coming up and down on this transponder, which explains the deviations in payload usage.

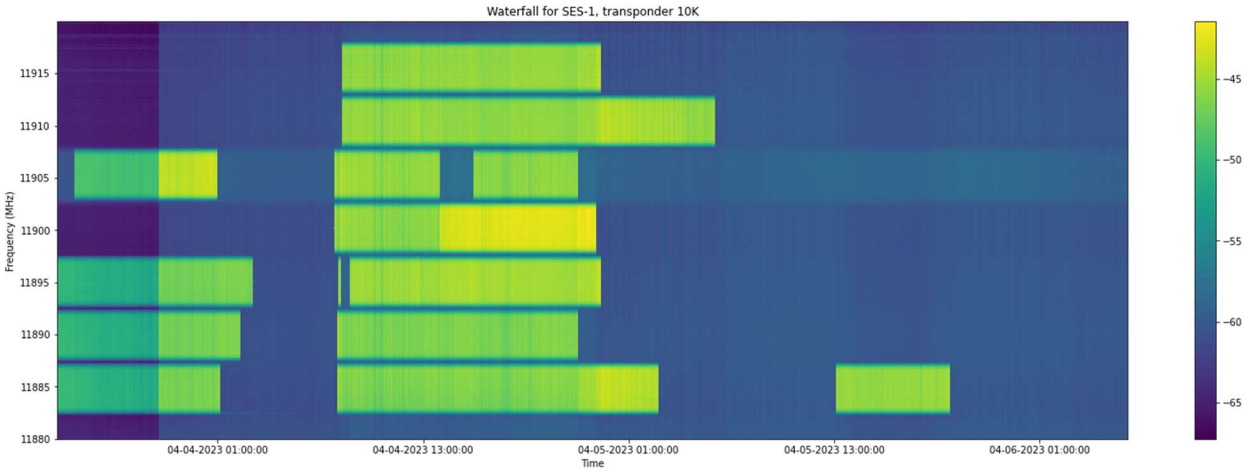

Fig. 19. Waterfall for SES-1, Transponder 10K

As shown in Fig. 20 and Fig. 21, we apply the ensemble model to the Symbol Rate, Interference, and Es/N0 to determine if anything out of ordinary.

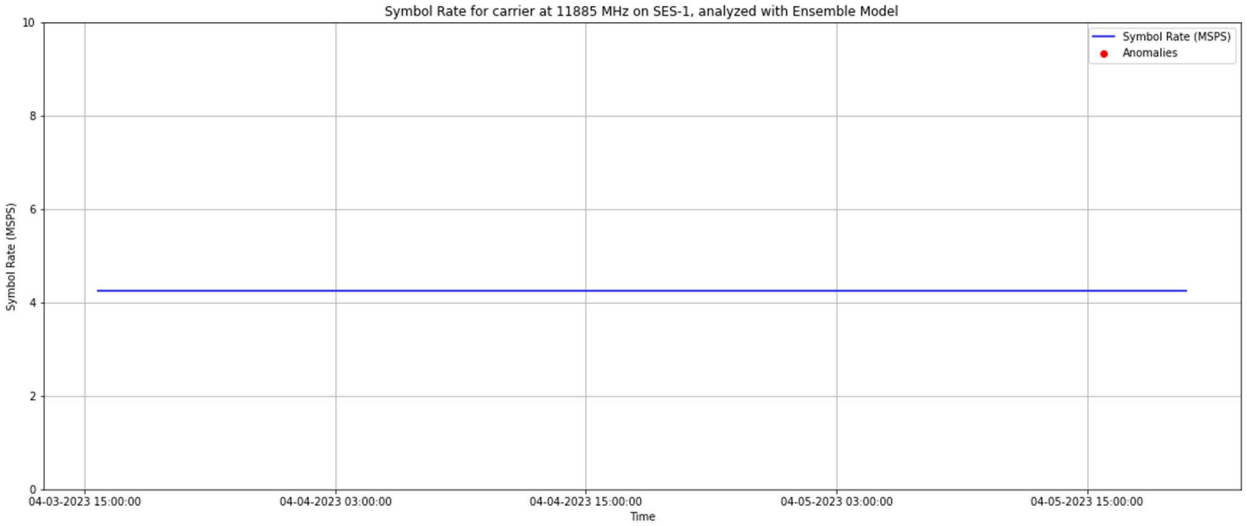

Fig. 20. Symbol Rate Analyzed with Ensemble Model for Carrier Centered at 11885 MHz on SES-1

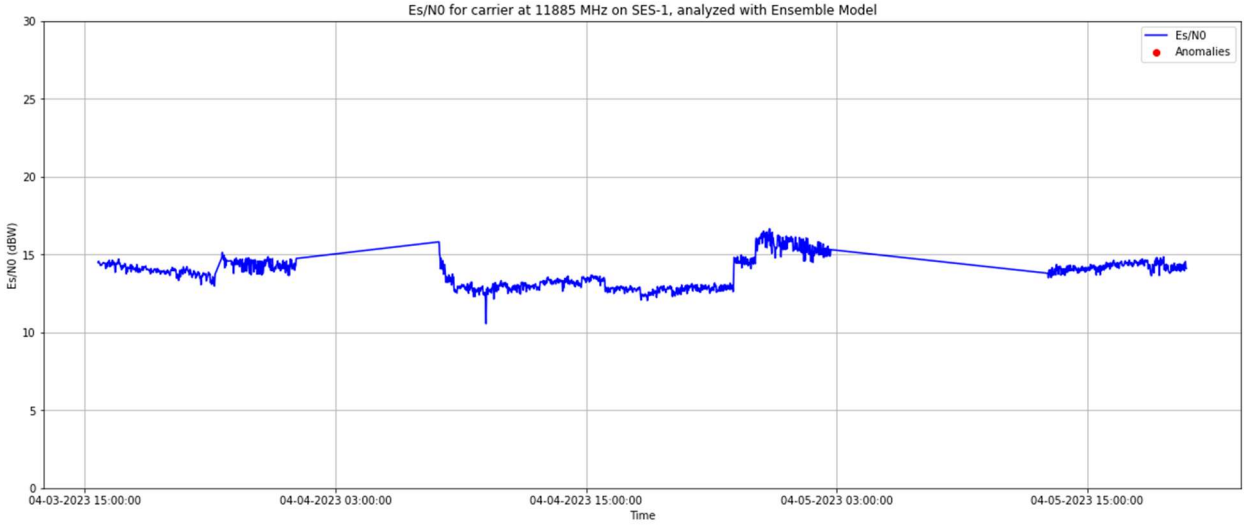

Fig. 21. Es/N0 for the SES-1 Carrier Centered at 11885 MHz

We then analyze the carrier that is up the most, the one centered at 11885 MHz. As shown in Fig. 22, none of the metrics on this carrier have any anomalies. As such, this might be attributed to normal day-to-day behavior on this transponder.

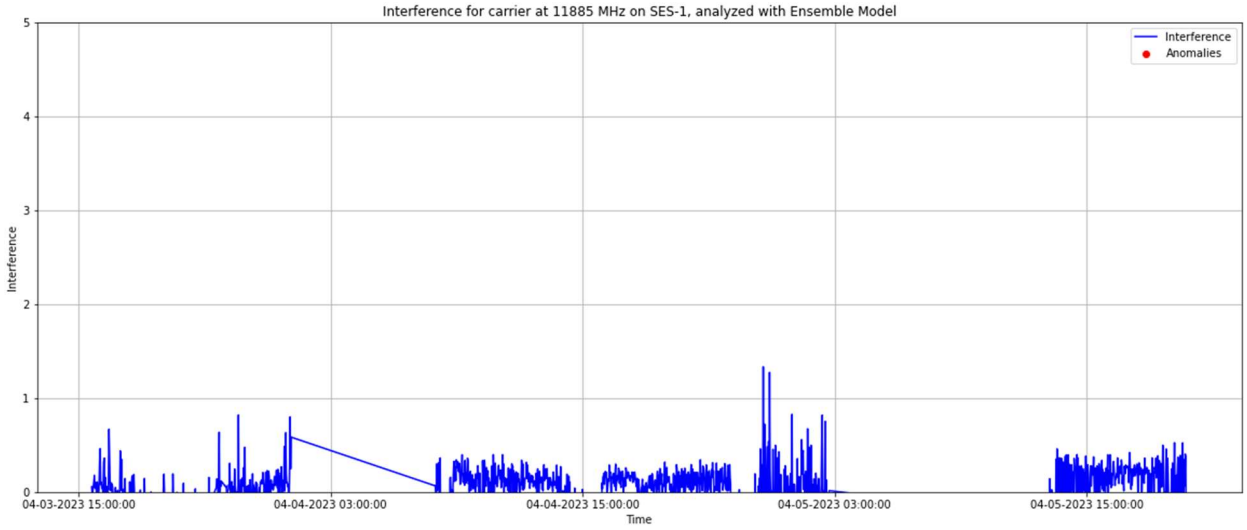

Fig. 22. Interference for the SES-1 Carrier Centered at 11885 MHz

As our final example, we will explore the ensemble model applied to Yamal 401, a Russian communications satellite. (Note that this is visualized with Sliding Gaussian + Moving Exponential Average for clarity.) In contrast to the previous examples, we start from the bottom up for this analysis. As shown in Fig. 23, anomalies are detected on the interference metric for a weekend's worth of RF transmissions collected from transponder 2 on the carrier centered at 11060 MHz.

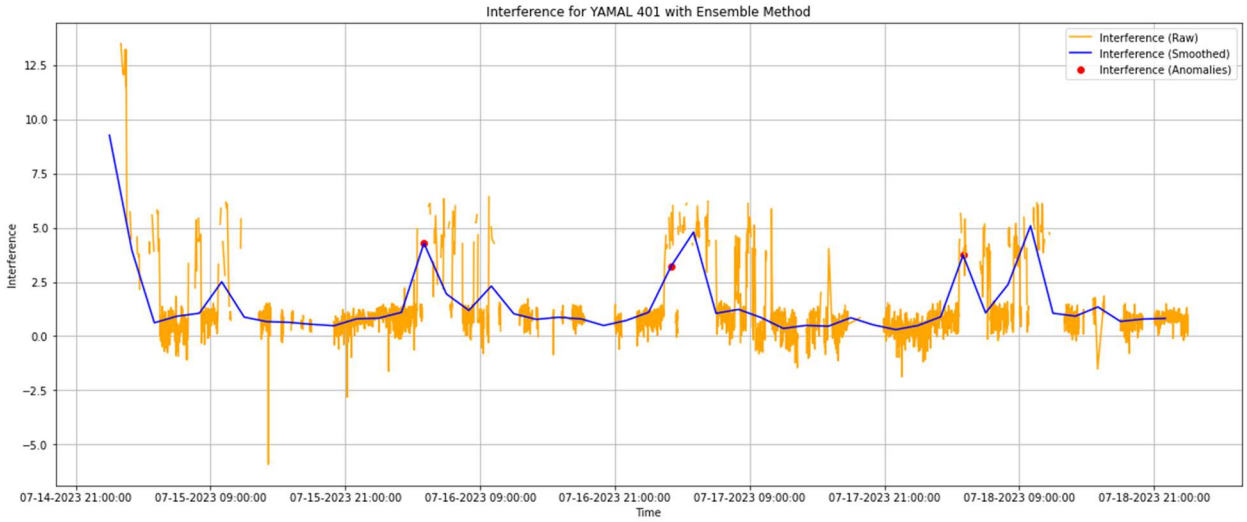

Fig. 23. Ensemble Model Applied to Yamal 401 Interference

Upon further investigation of the PSD, we see a sweeper signal. A sweeper is a narrow signal that sweeps continuously across frequency with the intent of degrading the signal quality. When visualized on a PSD, the sweeper will appear as a bump being transmitted on top of a carrier. This is an example of deliberate interference being applied by a malicious party.

Fig. 24 shows nominal usage at a time before the interference began.

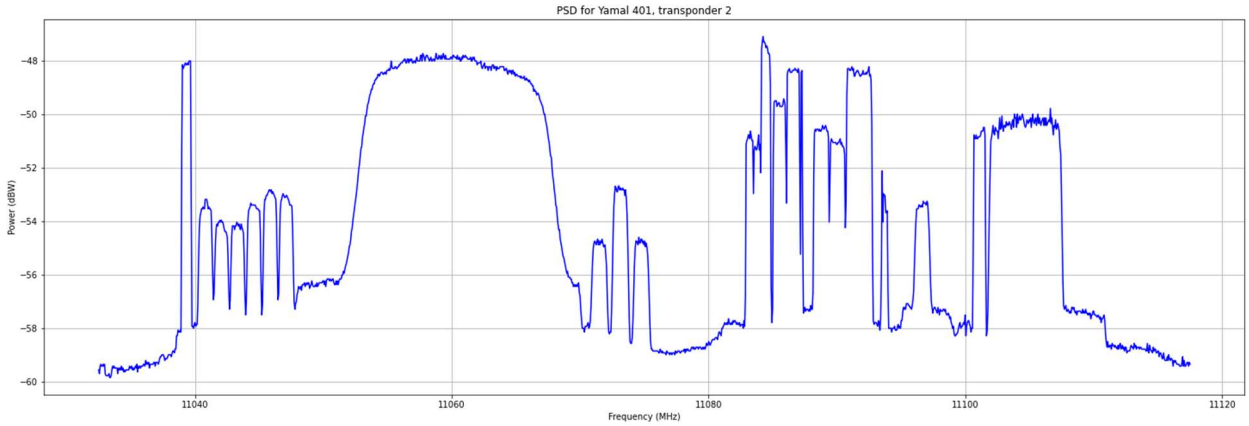

Fig. 24. Yamal 401 PSD

Fig. 25 shows interference on the carrier. (Note that this PSD's timestamp corresponds to the first anomaly detected by the ensemble model in Fig. 23.)

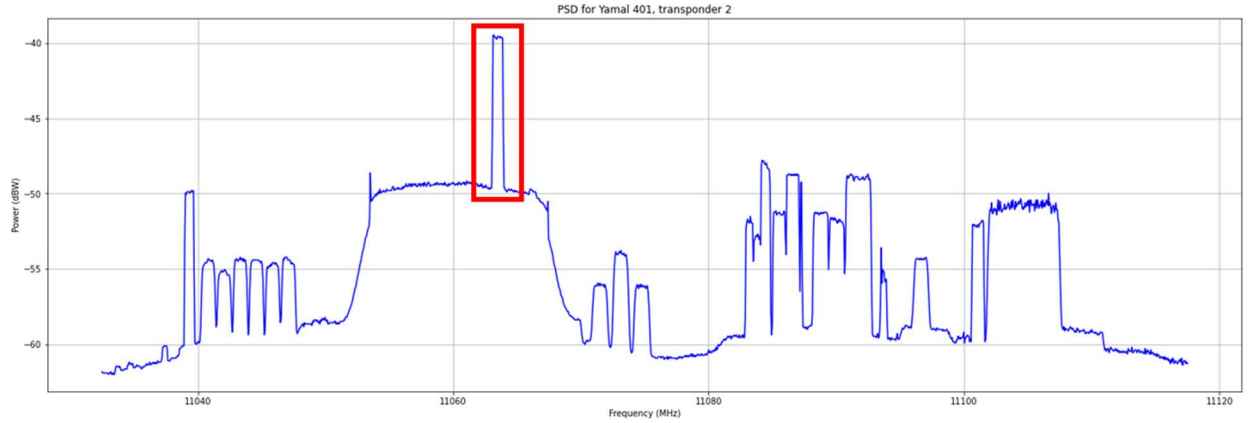

Fig. 25. Yamal 401 Trace with the Sweeper Highlighted in Red

Furthermore, this can be confirmed by analyzing the waterfall in Fig. 26. The yellow distortions highlighted in red represent the sweeper signal. Those familiar with RF interference might classify this as a hopper signal. However, the time between data samples indicate that a hopper might look identical to a sweeper on this waterfall.

This anomaly detection might lead an RF analyst to geolocate the sweeper signal. Given the current Russia/Ukraine war, it is highly likely that this interference could be from Ukraine or Ukrainian allies.

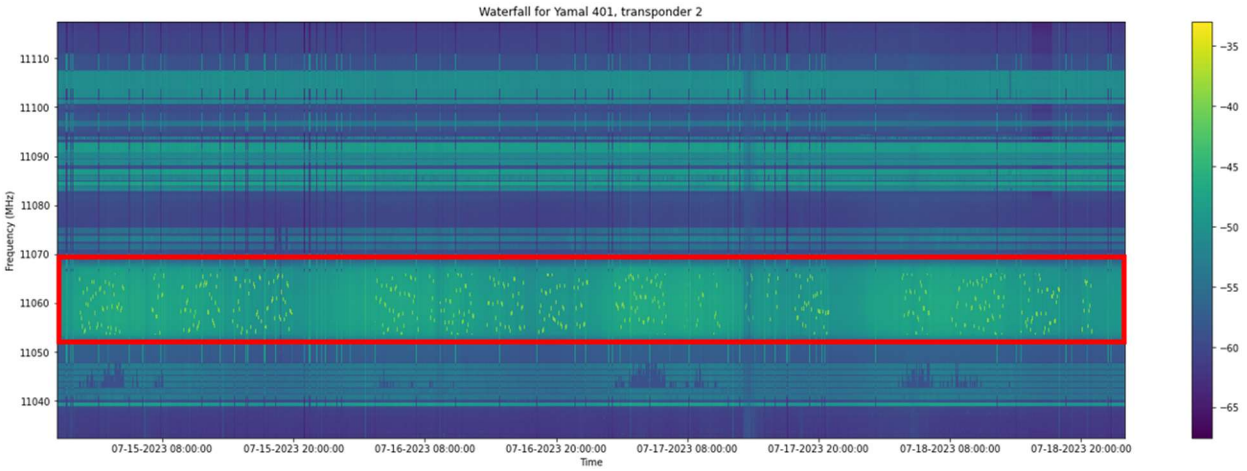

Fig. 26. Waterfall for Yamal 401, Transponder 2 with Interference Highlighted in Red

8. CONCLUSIONS

Based on the techniques evaluated in this paper, we recommend using the Sliding Gaussian + Exponential Weighted Moving Average, Ruptures, and Autoencoder + LSTM methods in an ensemble model for pattern-of-life analysis of RF data. This proposed ensemble model is lightweight; easy to customize at scale; and can be easily tuned as needed for individual transponders and carriers.

Additionally, RF analysts can mix and match these pattern-of-life analysis techniques to configure the most appropriate ensemble model. For instance, if an RF analyst wanted the ensemble model to be extremely sensitive to even slight changes in RF metrics, the Naïve approach might be preferred over the Exponential Weighted Moving Average approach.

Furthermore, since the techniques described in this paper are for time series datasets, the application is not limited to RF datasets. For example, the ensemble model could also be fed with telemetry data and orbital analysts could be alerted in the same way.

9. REFERENCES

[1] Charles Truong, Laurent Oudre, Nicolas Vayatis. “ruptures: change point detection in Python” Jan. 2, 2018. [Online] Available: https://arxiv.org/pdf/1801.00826.pdf.

[2] Duke, [Online]. Available: https://people.duke.edu/~rnau/411arim.htm [Accessed August 2023].

[3] R. Killick, P. Fearnhead, I. A. Eckley. “Optimal detection of changepoints with a linear computational cost” Oct. 9, 2012. [Online] Available: https://arxiv.org/abs/1101.1438.

[4] Sima Siami-Namini, Akbar Siami Namin. “Forecasting Economics and Financial Time Series: ARIMA vs. LSTM” March 15, 2018. [Online] Available: https://arxiv.org/ftp/arxiv/papers/1803/1803.06386.pdf.