Satellite Communication Security: Evaluation of Anomaly Detection Models

Clifa Mascarenhas[1], Dr. Nilesh B. Fal Dessai[2]

[1]*Student, Department of Information Technology and Engineering, Goa College of Engineering, Farmagudi, Goa, India*
[2]*Head of Department, Department of Information Technology and Engineering, Goa College of Engineering, Farmagudi, Goa, India*

---***---

Abstract - *Satellite-Based Communication (SATCOM) systems are increasingly pivotal for connecting remote areas, driven by advancements in manufacturing and radio technologies. However, this surge in adoption comes with heightened cybersecurity challenges. Neglected security services, coupled with evolving cyber threats, create an expanded attack surface for SATCOM.*
The study systematically compares the performance of diverse anomaly detection models on a synthetic dataset, assessing their ability to identify and counteract deviations from normal system behavior.

The outcomes of this research will contribute valuable insights to the selection of effective anomaly detection models for SATCOM security. Stakeholders involved in the SATCOM domain will benefit from this research, providing a foundation for informed decision-making when implementing security measures to safeguard critical communication channels.

Key Words: **SATCOM, cybersecurity, anti-jamming, integrity, confidentiality, anomaly detection**

1.INTRODUCTION

In an era marked by the imperative of global connectivity, Satellite-Based Communication (SATCOM) systems have emerged as pivotal facilitators, connecting remote regions and fostering seamless communication. As the demand for robust connectivity continues to rise, this project embarks on a comprehensive study, focusing on the intricate landscape of security threats, solutions, and challenges inherent in the deployment and operation of SATCOM systems, with a specific emphasis on anomaly detection.

Anomaly Detection Models: The primary objective of this project is to explore and evaluate various anomaly detection models tailored for SATCOM systems. Anomalies, deviations from the expected behavior, can signify potential security breaches or operational irregularities. This study delves into the effectiveness of state-of-the-art models, assessing their precision, recall, and overall performance in identifying anomalies within SATCOM data.

Security Threat Landscape: While the project narrows its focus to anomaly detection, it remains cognizant of the broader security threat landscape. Cyber threats, deliberate interference, and signal jamming are acknowledged as potential sources of anomalies, prompting the need for advanced detection mechanisms.

In essence, this project serves as a focused exploration into the application of anomaly detection models to the security dimensions of SATCOM systems. By bridging theoretical analyses with practical considerations, this endeavor aims to contribute valuable insights to the ongoing discourse on securing and sustaining SATCOM systems in an ever-evolving digital landscape.

2. Related Works

The paper[1] employs a robust ground truth methodology, labeling a specific number of segments in each telemetry variable to establish the accuracy of anomaly detection methods. Out of 21,644 samples, 287 are labeled as fake anomalies. Notably, the Deviation Divide Mean over Neighbors (DDMN) consistently achieves 100% precision and maintains a high F1-score above 90%. DDMN's F1-score reaches 0.969 when H = 2, surpassing DDM(next), DDM(prior), Z-SCORE, GMM, and K-means by significant margins of 93%, 92%, 30%, 61%, and 126%, respectively. Importantly, DDMN exhibits robustness to variations in the threshold (H). In contrast, other methods such as DDM(next), DDM(prior), and Z-SCORE exhibit low precision when H is small, as they tend to detect rapidly changing values as fake anomalies. GMM and K-means perform poorly in recall, erroneously assigning fake anomalies to the normal cluster.

This paper proposes a data-driven anomaly detection framework for satellite telemetry with fake anomalies. The authors propose the Deviation Divide Mean over Neighbors (DDMN) method to solve the fake anomaly problem caused by data errors in satellite telemetry data. They then use Long Short-Term Memory (LSTM), a deep learning method, to model multivariable time-series data and a Gaussian model to detect anomalies.

The authors applied their approach to telemetry data collected from sensors on an in-orbit satellite for more than two years and demonstrated its superiority. They also explored what conditions could lead to false alarms. The approach has been deployed to the ground station to monitor the health status of the in-orbit satellites.

Anomalies are events that differ from the usual behavior of system data, and analyzing satellite telemetry data is crucial in the development of aeronautics and astronautics. Satellites are complex systems composed of interrelated and mutually restricted components, and hardware monitoring is difficult. A single failure in a component or subsystem can be fatal to the system.

The Out-Of-Limit (OOL) method is widely used in current satellites, but it is not robust enough to detect various types of anomalies and requires significant domain knowledge and expertise from operators. To address these shortcomings, data-driven methods have been proposed. These methods use the stored telemetry to create a mathematical model of the nominal behavior of the satellite, as anomalous events are rare. Data-driven methods generate models from data and diagnosis directly from vast newly acquired telemetry, rather than building it based on human expertise.

In conclusion, this paper presents a data-driven anomaly detection framework for satellite telemetry with fake anomalies, addressing the limitations of current systems.

[2] This survey article is a comprehensive exploration of space anomaly detection, vital for ensuring space system safety amid increasing threats. It addresses specific challenges such as scalability, real-time detection, limited labeled data, concept drift, and adversarial attacks in space security. The article reviews current approaches and introduces an innovative integration of stream-based and graph-based methods for dynamic space anomaly detection, enhancing accuracy by capturing complex dependencies. It provides valuable insights and guidance for researchers and practitioners, emphasizing the need for innovative approaches to counter evolving space threats. The article concludes with future directions, offering a roadmap to advance anomaly detection capabilities and secure space networks.

It focuses on the importance of space anomaly detection in ensuring the safety and reliability of space systems, especially in the face of increasing threats. The article provides a thorough analysis and synthesis of state-of-the-art anomaly detection systems, identifying key challenges such as scalability, real-time detection, limited labeled data, concept drift, and adversarial attacks. By extensively reviewing existing approaches and methods, the article evaluates their strengths, limitations, and potential applications in space networks. It goes beyond a mere summary by introducing an innovative integration of stream-based and graph-based methods for dynamic space anomaly detection. This integration not only opens up new avenues for research but also enhances detection accuracy by capturing the complex temporal and structural dependencies within space networks.

The space industry is experiencing rapid growth and holds significant promise for the future, with a projected value exceeding $1 trillion by 2040. This expansion is driven by various factors, including decreasing launch costs and advancements in computing technology, particularly miniaturization. The emergence of small satellites and the availability of commercial off-the-shelf (COTS) components have made space more accessible to new entrants, resulting in increased satellite usage in communication, government operations, meteorology, remote sensing, and navigation.

The growth of the space industry has profound implications across different sectors, including job creation, technological advancements, increased accessibility, and potential economic growth. The expansion is expected to create numerous job opportunities in engineering, science, manufacturing, and support services. Advancements in technology for space missions can drive scientific research, improve communication systems, and enhance national security.

Space Information Networks (SINs) have gained significant attention within the space industry, encompassing a network of systems such as satellites, unmanned aerial vehicles, airships, and other objects capable of receiving, processing, and transmitting real-time spatial and temporal information. Satellite constellations offer tremendous potential in disaster response, scientific research, environmental management, and military missions.

3. Methodology

3.1 Dataset Overview

The dataset utilized in this study is a bespoke compilation generated to explore the intricacies of satellite telemetry data under varying environmental conditions. Simulated through a MATLAB script specifically designed for satellite communication scenarios, the dataset incorporates key variables such as temperature, humidity, signal strength, and time stamps. Comprising 50000 samples, the dataset adopts a time series format with irregular intervals, mirroring the unpredictability of real-world telemetry data. Each variable is recorded numerically, with temperature measured in degrees Celsius, humidity as a percentage, signal strength in decibels (dBm), and time stamps provided in UNIX format. The dataset is enriched by dynamic variations in environmental conditions, enabling a nuanced analysis of satellite communication scenarios. Prior to analysis, the data underwent preprocessing steps, including the normalization of temperature and humidity, along with the removal of outliers in signal strength. It is important to note that the dataset is intended for internal research purposes and serves as a valuable resource for investigating the impact of environmental factors on satellite communication reliability.

3.2 Models Used

The three models used in the anomaly detection and evaluation project are:

1. Isolation Forest:
Isolation Forest, introduced by Liu et al. in 2008[3], is an unsupervised machine learning algorithm designed specifically for anomaly detection. It operates on the principle of isolating anomalies by employing decision trees. Unlike traditional methods that rely on proximity-based metrics, Isolation Forest utilizes a fundamentally different approach by randomly selecting feature subsets and isolating anomalies efficiently in fewer splits.

The core mechanism of Isolation Forest lies in the construction of isolation trees. These trees are built recursively by randomly selecting a feature and then selecting a random split value within the range of the selected feature. This process continues until the instances are completely isolated, forming terminal nodes. Anomalies are expected to be isolated more quickly compared to normal instances due to their distinctiveness, resulting in shorter paths in the tree.

The average path length in the tree, denoted as E(h), is computed as follows:
E(h)=c(h)

Where:
c(h) is the average depth of a terminal node in a binary tree with
n
n instances.
The anomaly score for a data point x is calculated using the average path length and the number of isolation trees (T):

$(x, T)=2^{-E(h(x))/c(n)}$

Where:
E(h(x)) is the average path length for x across all trees,
c(n) is the average path length for n instances.

One of the distinguishing features of Isolation Forest is its ability to handle high-dimensional data effectively. Traditional methods often struggle with the curse of dimensionality, where the performance deteriorates as the number of features increases. However, Isolation Forest's random selection of feature subsets mitigates this issue, making it well-suited for datasets with a large number of dimensions.

Furthermore, Isolation Forest does not require prior knowledge or training on normal data, making it particularly useful in scenarios where labeled data is scarce or impractical to obtain. Its unsupervised nature allows for the detection of novel anomalies that may not conform to known patterns or distributions.

2. One-Class SVM (Support Vector Machine):
One-Class SVM is an unsupervised machine learning algorithm designed specifically for anomaly detection. Unlike traditional SVM, which is primarily used for binary classification tasks, One-Class SVM is trained only on instances considered normal during training. It aims to determine the hyperplane that best separates normal data from outliers, effectively creating a boundary around the normal instances in the feature space.
The core principle behind One-Class SVM is to find the optimal hyperplane that maximizes the margin around the normal data points while minimizing the penetration of outliers into this margin. Mathematically, this can be formulated as an optimization problem:
$\min_{w,,\rho} \frac{1}{2} \|w\|^2 +\frac{1}{vn} \sum_{i=1}^{n} \xi i{-}\rho$

Subject to:
$f(x_i)=\langle w, \phi(x_i\ )\rangle-\rho\leq \xi_i$
$\xi_I \geq 0$
$\rho \geq 0$

Where:
w is the weight vector,
ξ_I are slack variables representing the deviation of each training sample from the decision boundary,
ρ is the offset of the hyperplane from the origin,
ν is a parameter that controls the trade-off between maximizing the margin and controlling the fraction of outliers,
n is the number of normal instances in the training set,
$\phi(x_i)$ denotes the feature mapping function.
The decision function of the One-Class SVM can be expressed as:

$f(x) = \langle\, w, \phi(x)\rangle - \rho$

Where:
x is a new data point.

One of the key features of One-Class SVM is its ability to model the distribution of normal instances without relying on labeled data for anomalies. This makes it particularly effective in scenarios where labeled anomaly data is scarce or expensive to obtain.

Furthermore, One-Class SVM is well-suited for datasets with limited anomalies, where the majority of instances are considered normal. By focusing solely on the characteristics of normal data during training, One-Class SVM can learn a representation of normality that generalizes well to unseen data.

3. LSTM (Long Short-Term Memory):
LSTM is a variant of RNN designed to address the vanishing gradient problem, which impedes traditional RNNs from effectively learning long-range dependencies. It achieves this through a sophisticated gating mechanism that selectively retains or discards information over time, allowing for the preservation of relevant information across distant time steps. This mechanism consists of three gates: the input gate, the forget gate, and the output gate, each responsible for controlling the flow of information through the cell state.

The core mechanism of LSTM revolves around the manipulation of its internal state, known as the cell state, which serves as a conveyor belt for information flow across time steps. At each time step, the LSTM unit makes decisions on whether to update the cell state, erase certain information, or output specific information based on the input data and its current state. This process enables LSTM to capture and retain relevant temporal dependencies crucial for anomaly detection.

Mathematically, the operations within an LSTM cell can be summarized as follows:

$$f_t = \sigma\,(W_f \cdot [h_{t-1}, x_t] + b_f)$$
$$i_t = \sigma\,(W_i \cdot [h_{t-1}, x_t] + b_I)$$
$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$
$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$$
$$o_t = \sigma\,(W_o \cdot [h_{t-1}, x_t] + b_o)$$
$$h_t = o_t * \tanh(C_t)$$

Where:
f_t is the forget gate output,
i_t is the input gate output,
$\tilde{C}_t$ is the candidate cell state,
C_t is the cell state,
o_t is the output gate output,
h_t is the hidden state output,
$W_{f,i}, W_C, W_o$ are weight matrices for each gate,
b_f, b_i, b_C, b_o are bias terms for each gate,
x_t is the input at time step
h_{t-1} is the hidden state at the previous time step.

LSTM's ability to capture long-term dependencies and intricate temporal patterns makes it particularly effective for anomaly detection in multivariate time-series data. Traditional methods often struggle to model complex relationships across time, especially in the presence of noisy or irregular data. LSTM's hierarchical structure and sequential learning capability empower it to discern subtle deviations from normal behavior, thereby facilitating the detection of anomalies.

Furthermore, LSTM's adaptability to varying sequence lengths and its robustness to noisy inputs enhance its applicability across diverse domains. Whether it's monitoring industrial processes, detecting anomalies in network traffic, or predicting financial market trends, LSTM's versatility and efficacy make it a go-to choice for time-series analysis tasks.

These models were selected to provide a diverse set of approaches for anomaly detection, ranging from traditional machine learning algorithms like Isolation Forest and One-Class SVM to a deep learning approach with LSTM. Each model has its strengths and limitations, and the evaluation metrics were employed to compare their performance on the specific dataset.

3.3 Flow of the project

1) Dataset Preparation:

Create and label a dataset representative of SATCOM communication scenarios. Include variations in signal-to-noise ratios, frequencies, and transmission power levels.

2) Model Selection and Training:

Explore diverse anomaly detection models such as Isolation Forest, One-Class SVM, and LSTM. Train and evaluate each model on the prepared dataset to assess their performance in detecting anomalies.

3) Evaluation Metrics:

Employ standard evaluation metrics like precision, recall, and F1-score to quantify the effectiveness of anomaly detection models. Compare and analyze the performance of different models to identify the most suitable for SATCOM security.

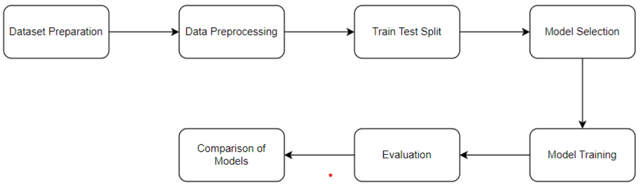

Fig 1. Flow diagram

3.4 Equations used

1. Precision (Positive Predictive Value):

Precision is the ratio of true positive predictions to the total predicted positives. It measures the accuracy of positive predictions made by the model.

$$\text{Precision}=\frac{\text{True Positives}}{\text{True Positives + False Positives}}$$

2. Recall (Sensitivity or True Positive Rate):

Recall is the ratio of true positive predictions to the total actual positives. It measures the ability of the model to capture all the relevant instances.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives + False Negatives}}$$

3. F1 Score:

The F1 score is the harmonic mean of precision and recall. It provides a balance between precision and recall and is particularly useful when there is an imbalance between classes.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision + Recall}}$$

These metrics are commonly used in classification tasks to evaluate the performance of machine learning models, particularly in scenarios where imbalances between positive and negative classes exist.

4. RESULTS

The dataset is created in matlab and has 50000 samples 40000 normal samples and 10000 anomalous samples. The dataset consists of 10 features. All the models were tested using the dataset and the results are shown in the table below:

Comparison			
Metric	Isolation Forest	One-Class SVM	LSTM
Precision	0.11388177668930201	0.17356475300400534	1.0
Recall	0.5079443892750745	0.7745779543197616	1.0
F1-Score	0.18605074111121214	0.28358480276313397	1.0

Table -1: Comparison of Isolation Forest, One-Class SVM, and LSTM Models

Precision:

Isolation Forest: The precision for Isolation Forest is relatively low (0.11), indicating a high number of false positives. The model's predictions of anomalies have a high chance of including normal instances.

One-Class SVM: Similar to Isolation Forest, One-Class SVM also has a low precision (0.17), suggesting a notable number of false positives.

LSTM: The LSTM model shows a precision of 1.0, which means that all instances predicted as anomalies are indeed anomalies. There are no false positives.

Recall:

Isolation Forest: The recall for Isolation Forest is moderate (0.51), indicating that the model captures around half of the actual anomalies.

One-Class SVM: One-Class SVM has a higher recall (0.77) compared to Isolation Forest, suggesting a better ability to identify anomalies.

LSTM: The LSTM model shows a recall of 1.0, indicating that it captures all actual anomalies. There are no false negatives.

F1-Score:

Isolation Forest: The F1-score for Isolation Forest is low (0.19), reflecting a balance between precision and recall, but both are relatively low.

One-Class SVM: The F1-score for One-Class SVM is also low (0.28), suggesting a trade-off between precision and recall.

LSTM: The LSTM model achieves a perfect F1-score of 1.0, indicating a harmonious balance between precision and recall. This is because precision and recall are both 1.0.

In summary, the LSTM model appears to outperform Isolation Forest and One-Class SVM in terms of precision, recall, and F1-score. It achieves perfect scores, indicating that it correctly identifies anomalies without any false positives or false negatives.

5. CONCLUSIONS

Anomaly detection, a pivotal aspect of SATCOM security, was approached through the lens of three distinct models: Isolation Forest, One-Class SVM, and Long Short-Term Memory (LSTM) neural networks.

- Isolation Forest and One-Class SVM: Both models demonstrated varying degrees of success in identifying anomalies. Isolation Forest and One-Class SVM exhibited limitations in achieving a balance between precision and recall. Isolation Forest, while showing moderate recall, suffered from a high number of false positives, impacting precision. One-Class SVM demonstrated improved recall but still faced challenges in precision.

- LSTM: The LSTM model showcased exceptional performance, achieving perfect precision, recall, and F1-score. The LSTM model, with its deep learning capabilities, demonstrated a profound ability to discern anomalies without any false positives or false negatives.

The anomaly detection model comparison underscores the superiority of the LSTM model, indicating its potential for robust anomaly identification in SATCOM systems. However, the choice of the most suitable model depends on specific use cases, data characteristics, and deployment scenarios.

6. FUTURE SCOPE

The dataset used was generated in MATLAB and the results might vary for actual datasets. Given that the dataset was artificially generated in MATLAB, it's understandable that the LSTM model achieved perfect results.

In addition to the current findings, there exists a rich landscape of potential future directions to further enhance the efficacy and applicability of anomaly detection in satellite communication.

One avenue for future exploration involves the evaluation of the developed models on real-world datasets sourced from satellite communication systems, providing crucial insights into their performance under authentic operational conditions.

Furthermore, the incorporation of advanced feature engineering techniques tailored to satellite communication data could yield improvements in model robustness and detection capabilities.

Additionally, the optimization of model parameters and architectures, guided by domain-specific knowledge, presents an opportunity to refine and tailor the anomaly detection process to the intricacies of satellite communication systems.

Moreover, the exploration of dynamic adaptation and continual learning mechanisms could enable the models to evolve and adapt to changing communication patterns and emerging anomalies over time.

Cross-domain application of the developed models beyond satellite communication, such as in network security or financial fraud detection, offers avenues for broader impact and validation of their generalization capabilities.

Finally, considerations of scalability and efficiency are paramount for the deployment of anomaly detection models in large-scale satellite communication infrastructures, necessitating optimizations for real-time processing and resource-constrained environments.

By addressing these future research directions, we aim to propel the field of anomaly detection in satellite communication towards greater efficacy and practical utility.

REFERENCES

[1] Yakun Wang, Jianglei Gong, Jie Zhang, Xiaodong Han, "A Deep Learning Anomaly Detection Framework for Satellite Telemetry with Fake Anomalies", International Journal of Aerospace Engineering, vol. 2022, Article ID 1676933, 9 pages, 2022. https://doi.org/10.1155/2022/1676933

[2] Abebe Diro, Shahriar Kaisar, Athanasios V. Vasilakos, Adnan Anwar, Araz Nasirian, Gaddisa Olani, “Anomaly detection for space information networks: A survey of challenges, techniques, and future directions”, Computers & Security, Volume 139, 2024, 103705, ISSN 0167-4048, https://doi.org/10.1016/j.cose.2024.103705.

[3] F. T. Liu, K. M. Ting and Z. -H. Zhou, "Isolation Forest," 2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy, 2008, pp. 413-422, doi: 10.1109/ICDM.2008.17.

[4] N. Abdelsalam, S. Al-Kuwari and A. Erbad, “Physical layer security in satellite communication: State-of-the-art and open problems”, 2023, [online] Available: https://arxiv.org/abs/2301.03672.

[5] Zixiang Jia, “Anti-jamming Technology in Small Satellite Communication” 2018 J. Phys.: Conf. Ser. 960 012013

[6] Dan Shen, Gang Wang, Genshe Chen, Khanh D. Pham, Erik Blasch, “Network survivability oriented Markov games (NSOMG) in wideband satellite communications” 2014, 2014 IEEE/AIAA 33rd Digital Avionics Systems Conference (DASC)

[7] Xi, Y.; Liu, J.; Zhao, W. “SATCOM Earth Station Arrays Anti-Jamming Based on MVDR Algorithm”. Appl. Sci. 2023, 13, 8337.

[8] Li, Jiong. (2021). Satellite communication anti-jamming based on ABC blind source separation. 10.21203/rs.3.rs-279435/v1.

[9] I. Ahmad, J. Suomalainen, P. Porambage, A. Gurtov, J. Huusko and M. Höyhtyä, "Security of Satellite-Terrestrial Communications: Challenges and Potential Solutions," in IEEE Access, vol. 10, pp. 96038-96052, 2022, doi: 10.1109/ACCESS.2022