

AI Generated Wireless Data for Enhanced Satellite Device Fingerprinting

Ningning Wang[†], Tianya Zhao[†], Shiwen Mao[‡], Xuyu Wang[†]

[†]Knight Foundation School of Computing and Information Sciences, Florida International University, Miami, FL 33199

[‡]Department of Electrical and Computer Engineering, Auburn University, Auburn, AL 36849, USA

Emails: nwang012@fiu.edu, tzhao010@fiu.edu, smao@ieee.org, xuywang@fiu.edu

Abstract—While Low Earth Orbit (LEO) satellite communications have attracted more and more attention recently, its security issues for wireless communications become greatly important. Implementing Radio Frequency (RF) fingerprinting by leveraging In-phase/Quadrature-phase (I/Q) satellite data is a reliable way to enhance wireless communication security. However, this method still faces two challenges. The first is the high cost of acquiring I/Q data from LEO satellites, and the second is the wide variation of data collected in different periods. To solve these two challenges, we develop a modified Variational Autoencoder (VAE) with a tailored loss function to generate LEO satellite data. Then, to achieve a multi-period model generalization, a cosine similarity classifier is incorporated into prototypical networks for facilitating few-shot learning at LEO satellite data in different periods. In our experiments, the proposed model can achieve a 50-classification accuracy of 99.80% in a single period with virtual satellite data generated by the modified VAE, which is higher than the classification result of 98.05% using only real data. Besides, using few-shot learning, our results demonstrate the effectiveness of model adaptation, obtaining a significant improvement in classification accuracy, from 2% to 77% of the original model.

Index Terms—RF Fingerprinting, Data Generation, Few-shot Learning, LEO Satellite

I. INTRODUCTION

The constant exposure of wireless devices to various threats, particularly impersonation attacks, has become a widespread security concern [1]. To address the security issue of wireless devices, Radio Frequency (RF) fingerprinting techniques are proposed, which leverage electromagnetic waves emitted by wireless devices to identify distinctive hardware defects resulting from manufacturing imperfections [2]. This method has a significant advantage, as it is difficult to replicate with other wireless devices. RF fingerprinting techniques have been exploited for different wireless techniques such as Bluetooth, WiFi, and RFID devices. Recently, deep learning approaches have been used to process and analyze in-phase (I) and quadrature (Q) signal data from wireless devices, achieving better results in multi-device identification scenarios involving 10,000 devices [3]. Thus, RF fingerprinting can be leveraged as a robust countermeasure against impersonation attacks [4] [5] [6], providing its potential applicability across various wireless communication technologies.

Low Earth Orbit (LEO) satellites have the potential to reduce transmission delay, mitigate path loss, and establish global coverage through constellations comprising multiple

satellites [7]. This configuration holds promise for frequency reuse, presenting a prospective paradigm for satellite-based mobile communications. However, LEO satellite systems face security risks, including spoofing and replay attacks. Cryptographic methods have some limitations, while commonly used for traditional authentications. For example, legacy satellites do not adopt cryptographic techniques due to inherent constraints like processing power and high cost. Additionally, cryptographic approaches may be vulnerable to attacks. As the development of LEO satellite technology continues, there arises a critical need for the consideration of the physical layer security of these systems [8]. Thus, the application of RF fingerprinting for the authentication of LEO satellite devices appears as an effective and reliable method.

However, the I/Q data collection from LEO satellites poses considerable challenges. The elevated orbital altitudes and swift velocities of these satellites relative to Earth result in a limited capture window for each satellite within a given period. The temporal constraint limits that a particular satellite remains within the antenna's capture range for no more than one hour. In a recent study by Smailes et al. [9], over 1.7 million Iridium data points were obtained through high-performance software-defined radio (SDR). However, the collected data exhibits suboptimal performance, with a classification accuracy of approximately 60%. In a different approach in [10], I/Q data over two months was collected using a USRP X310 device, thus achieving an accuracy exceeding 80% through a ResNet18 network [11]. In fact, the high cost of satellite data collection becomes a main bottleneck for RF fingerprinting. To address the issue, data augmentation schemes can be exploited. For example, Ding et al. [12] implemented a data augmentation strategy for low-resource tagging tasks. Tomczak and Welling proposed a competitive Variational Autoencoder (VAE) model, demonstrating efficacy in comparison to Generative Adversarial Networks (GAN) models [13] [14]. Mansi et al. [15] leveraged conditional GAN for wireless modulation data augmentation to improve the classification performance. Thus, the integration of Artificial Intelligence (AI) generative modeling methodologies with satellite data emerges as a promising avenue for further exploration in this paper.

Moreover, because of the divergent operational velocities and orbital characteristics in LEO satellite systems, data acquisition at the same geographic coordinates during different

time periods may lead to significant variations. Consequently, addressing the challenge of enhancing model generalization across distinct scenarios warrants thoughtful consideration. Notably, Snell et al. [16] have underscored the efficacy of few-shot learning methodologies in augmenting the performance of models operating in the spectrum domain. Steven et al. [17] exploited prototypical networks for cross-domain Wi-Fi device fingerprinting over different transmission distances, and different times.

In this paper, we develop a modified VAE model for LEO satellite data generation, thus mitigating the cost of LEO data collection and improving the LEO device classification performance. In addition, we also integrate the synthetically generated data with prototypical networks to enhance the generalization capabilities of the model over different domains such as different time periods. This utilization of generated LEO satellite data can effectively improve the classification performance of models initially trained at the same time, as well as enhance the generalization capacity of the model at a low cost using few-shot learning.

In summary, the main contributions in this paper include:

- 1) To the best of our knowledge, this is the first work that harnesses limited satellite data to generate new I/Q data and achieve data augmentation. The generated data successfully helps the learning model to improve the accuracy of recognizing LEO satellites.
- 2) For the generated data, we utilize the Frechet Inception Distance (FID) to measure its diversity and similarity to the original data. Also, we develop prototypical networks to improve the generalization ability of the model to cope with the cross-domain problem of satellite signals over different time periods.
- 3) We experimentally evaluate the effectiveness of our proposed generated method. The proposed model can achieve a 50-classification accuracy of 99.80% in a single period by the modified VAE. Besides, by using few-shot learning, our approach can obtain a large improvement in classification accuracy, from 2% to 77% of the original model.

The remainder of this paper is organized as follows. Section II is the preliminaries and motivation and Section III is the system design with satellite data. The experimental study is in Section IV. Section V concludes this paper.

II. PRELIMINARIES AND MOTIVATION

In this section, we will introduce the I/Q data of LEO satellites, especially on data offset due to time effect. Then, we show how to introduce the generated data and transfer learning to overcome these problems.

A. RF Fingerprinting of Satellite

In light of the inherent introduction of defects during the manufacturing process of hardware devices, the signals manifest discernible characteristics within the physical layer. This characteristic is transmitted in wireless communication




Fig. 1: Transmitted signal represented by I/Q data.

of the device along with the transmission of the in-phase (I) and quadrature (Q) signal. Fig. 2 shows the modulation, propagation, and demodulation processes of I/Q signals. $s(t)$ is the RF signal which is created with the $I(t)$ and $Q(t)$ baseband signals and it can be formulated as follows [18],

$$s(t) = I(t) \cdot \cos(2\pi f_c t) - Q(t) \cdot \sin(2\pi f_c t), \quad (1)$$

where $s(t)$ is the corresponding signal and it can be split into $I(t)$ and $Q(t)$ centered around the carrier frequency f_c . In the demodulation process, the $I(t)$ and $Q(t)$ baseband signals can be recovered by mixing the received signal $y(t)$ with a local oscillator.

Generally, satellite data transmission is facilitated through Quadrature Phase Shift Keying (QPSK) modulation, a prevalent modulation scheme in wireless communication systems. The selection of QPSK modulation is of particular importance as it plays a crucial role in shaping the distinctive features exhibited by the transmitted signals. Consequently, this modulation choice contributes significantly to the delineation and analysis of physical layer attributes within the context of the used LEO satellite I/Q dataset.

TABLE I: The difference between LEO satellite datasets in two different periods

Label	Mean of data in P1	Mean of data in P2	Binary Classification Accuracy (%)	Adaptation Classification Accuracy (%)
0	0.5116	0.4712	98.56%	2.05%
1	0.4605	0.5083		
2	0.5445	0.5023		
3	0.4939	0.5082		
⋮	⋮	⋮		
49	0.496	0.5255		

Although LEO satellites have the advantages of signal robustness and minimal delay, their operating speed and orbital dynamics have a significant impact on the verification process at the physical layer. It is worth noting that the time of the dataset is split into two distinct parts, and the main difference between these parts deserves careful consideration. In order to intuitively reflect the differences between the data of different periods, we calculate the mean value of all I/Q data of the same label (i.e., the same satellite and antenna) of the two periods' data (P1& P2) and compute the cosine similarity between them.




Fig. 2: System block diagram.

The results are shown in Table I. We find that the cosine similarity of above 90% between two datasets (i.e., P1 and P2) cannot reflect the distribution difference. Thus, we use a binary classifier to verify the difference in the data between two periods. If there is no difference in the data, the classification accuracy should be ideal 50%. As shown in the fourth column of Table I, an accuracy of 98.56% has been achieved ultimately, which illustrates a large difference between P1 and P2 datasets. On the other hand, we use satellite I/Q data to train a Residual Neural Network (ResNet) model in the P1 data and test the performance in the P2 data. We find that the adaptation classification accuracy is only 2% in the P2 data, which further validates the distribution shift issue in the P1 and P2 datasets. Therefore, we can conduct data generation experiments within the same time period. Based on the above experimental results, we consider performing data generation on the P1 data and subsequently applying few-shot learning to the P2 data.

III. SYSTEM DESIGN WITH SATELLITE DATA

This section presents an overview of the comprehensive system architecture, as depicted in Fig. 2, consisting of two important components. The entirety of the system includes three discrete phases, each contributing to the experimental framework. First, a classification experiment is executed employing a conventional method [17]. Subsequently, a data augmentation procedure is initiated by utilizing a modified VAE model. This iterative process involves the generation of virtual LEO satellite data through the VAE model, which is then combined with the authentic dataset. Second, the augmented dataset is harnessed to refine the original model's performance. This process involves the integration of both virtual and real data, with the objective of enhancing the model's classification performance. Third, the adapted model is deployed on datasets originating from different temporal periods. A few-shot learning approach with prototypical networks is used to update the original model, facilitating improved adaptability and performance across diverse temporal contexts. This multi-phased system framework includes a holistic strategy aimed

at refining and extending the classification capabilities of the overall model.

A. Data Generation

VAE model is a type of generative model in which the distribution of the encoding is regularized during training. In other words, VAE provides a latent space structure. The internal structure of the multidimensional latent space for a well-learned model defines its properties. The decoder component reconstructs the input using this information. The core idea of VAE is to compress a random vector x in a high-dimensional space into a latent variable z in a low-dimensional space by variational encoding. This process can be simplified and expressed as follows,

$$P_\theta(x, z) = P_\theta(x|z)P_\theta(z), \quad (2)$$

where $P_\theta(z)$ denotes the prior distribution of the latent variable z , which is generally set to standard Gaussian distribution; $P_\theta(x|z)$ denotes the conditional probability density function of the input variable x when z is known, θ denotes the variable parameter. The posterior distribution $P_\theta(z|x)$ is intractable since the parameters θ and the latent variable z are unknown. The standard VAE uses a recognition model $q_\phi(z|x)$ as an approximation to the true posterior $P_\theta(z|x)$. VAE generates two outputs in the encoder including a mean vector and a standard deviation vector. These two vectors form the parameters of the latent variable z . This mechanism allows the encoder to learn different mean values corresponding to each potential class, and the standard deviation reduces the overlap of feature classes. The use of the univariate Gaussian distribution is a popular option in VAE. The Kullback–Leibler divergence measure between the approximation output and the target (input) characterizes the regularized term. Due to introducing a normal distribution, the loss function has been modified accordingly. The loss function can be expressed by,

$$loss_r = \|x - \hat{x}\|_2 = \|x - d_\theta(\mu_x + \sigma_x \epsilon)\|_2, \quad (3)$$

$$\mu_x, \sigma_x = e_\theta(x), \epsilon \sim N(0, I), \quad (4)$$

$$loss_k = -D_{KL}(N(\mu_x, \sigma_x) || N(0, I)), \quad (5)$$

$$loss = \alpha loss_r + \beta loss_k, \quad (6)$$

where x is the real value and $\hat{x}$ is the predicted value, μ_x and σ_x are the mean value and standard variance of value x , respectively. $loss_r$ is the reconstruction loss, which is the L2 norm. It is used to reduce the randomness of generated data. μ_x and σ_x are calculated by encoder $e_\theta(x)$. ϵ represents sampling a random noise term from a standard normal distribution. $loss_k$ is similarity loss and represents the KL divergence between the encoded distribution in the latent space and the standard normal distribution, where KL divergence is a method used to measure the difference between two probability distributions. Here, it is used to ensure that the distribution of the potential space is as close as possible to the standard normal distribution. And α and β denote the weights of $loss_r$ and $loss_k$, respectively. In this paper, we have modified the

weighting of the L2 norm and KL divergence components within the loss function, aiming to minimize the discrepancy between the generated and original values.

Since the I/Q signals are two mutually orthogonal components, the encoder of this VAE model decomposes the I/Q signals. The LEO signals will be compressed into latent distribution and then sampled to make the decoder reconstruct the signal based on the feature of the latent space. We can learn the features of satellite I/Q data by ResNet18. Although the I/Q data is received in temporal order, it is finally presented without temporal information, because the ResNet model is insensitive to temporal information for data feature extraction. In the decoder network, the deconvolutional network is used to reconstruct the original signal.

The loss function is composed of reconstruction loss and KL divergence. To reconstruct a better signal, the reconstruction loss measures the sum of all the squared differences between the output and the input. The KL divergence loss regularizes the latent distribution according to a standard Gaussian prior.

The FID serves as a metric to quantify the feature dissimilarity between authentic and generated datasets. Additionally, FID is employed as an evaluative criterion to assess the quality of the generated data. The following formula facilitates the computation of FID, that is,

$$FID(x, y) = \|\mu_x - \mu_y\| + Tr\left(\sum_x + \sum_y - 2\sqrt{\sum_x \sum_y}\right), \quad (7)$$

where x and y are the generated data and original data, respectively; μ_x and μ_y are the mean values of x and y , respectively; $\sum X$ is the covariance matrix, and Tr means the trace of the two matrices.

In this paper, we modify the loss function by reducing the weight of KL divergence while increasing the weight of regularization. The impact of this change can be reflected in the change in the value of the FID.

B. Model Adaptation

In the third module of our system, we employ prototypical networks for the implementation of model adaptation. The fundamental idea underlying prototypical networks is straightforward. For classification tasks, the network identifies the prototype center of each class within the semantic space. In few-shot learning, the prototypical network is trained to learn how to effectively fit these centers, thereby achieving a metric function capable of discerning its class within the prototype center of the metric space using a limited number of samples. Assume the original dataset is D . For each episode, it contains a support set and a query set, which can be expressed by,

$$D_{episode} = D_{support} \cup D_{query} = \{s_i\}_{i=1}^{n_s} \cup \{q_i\}_{i=1}^{n_q}, \quad (8)$$

where s_i is the i_{th} data in the support set and q_i is the i_{th} data in the query set. n_s and n_q are the sizes of the support set and query set, respectively. We have a small dataset of episodes by randomly selecting N categories in the original dataset D ,

including K data per category to form the support set, and Q data to create the query set. By analogy, we construct *episodes* of small datasets. For each data, feature extraction is performed using an Encoder as the following,

$$h_{s_i} = f_\varphi(s_i), \quad (9)$$

$$h_{q_i} = f_\varphi(q_i), \quad (10)$$

where ResNet18 is used as the Encoder $f_\varphi(\cdot)$; h_{s_i} and h_{q_i} are the embedding of each data. Then, we calculate the prototype for each category in the support set. The embedding table can be formulated by,

$$p_{c_j} = \sum_{\{i|l_{s_i}=c_j\}} h_{s_i}, \quad (11)$$

where l_{s_i} indicates the label of data s_i , c_j is the category of the j_{th} support set.

From Fig. 2, the dataset will be divided into the support set and the query set. All samples of a single task in the support set will be passed through an encoder $f_\varphi(\cdot)$ in this experiment to obtain the embedding points of the whole support set. The embeddings form an embedding table to represent the class, and the embeddings obtained from the query set are measured. Here, we use cosine similarity to calculate the distance between different classes.

During testing, the support set samples are utilized to compute the cluster center of the new class and a nearest neighbor classifier approach is employed for prediction. It is noteworthy that, in this work, the Euclidean distance is not used as the distance metric. Instead, we choose the cosine similarity as the metric function, aiming to capture the relative positional relationships between I/Q data points. This selection stems from the uncertainty regarding whether time-affected data will remain as pertinent as previous data. The support set samples are leveraged for feature extraction, with the mean of all samples in each class serving as the fundamental prototype vector. Then, the distances to each class are computed, and normalized to obtain probabilities, and stochastic gradient descent is applied to minimize this value. The cosine similarity classifier facilitates the learning of an embedding space, wherein similar features from the same class are closely clustered together, thereby enhancing the discriminative capacity of the feature extractor.

In the inference phase, the class mean is employed as the basic prototype for the new class, and the cosine similarity between the prototype and the sample is computed for classification. Our objective is to maximize the expected cosine similarity, as it is positively correlated with classification accuracy, which can be formulated by [19],

$$\begin{aligned} \text{Max}[E_P[E_X[\text{COS}(P, X)]]] &= \text{Max}[E_{P,X}[\bar{P} \cdot \bar{X}]] \\ &= \text{Max}[E_X[\bar{X}] \cdot E[\frac{P}{\|P\|_2}]], \end{aligned} \quad (12)$$

where $\text{COS}(\cdot, \cdot)$ is the cosine similarity function, X is the feature of a class, P means the prototype, $\bar{X}$ and $\bar{P}$ are the normalized feature and normalized prototype, respectively.




Fig. 3: Comparison between the (a)real data and (b)generated data.

TABLE II: FID corresponding to different loss function weights.

$\alpha \backslash \beta$	0.25	1	1.5
0.25	5.502	5.487	5.630
1	3.794	5.514	5.810
1.5	3.280	5.271	5.540

$E_X[\bar{X}]$ means the expected value of the normalized feature $\bar{X}$. $E[\frac{P}{\|P\|_2}]$ means the expected value of the unit vector pointing in the normalized prototype $\bar{P}$.

In the context of model adaptation, we explore the utility of generated I/Q data. The experimental configuration employs a 50-way 5-shot design, where the support set is comprised of 5 samples per label from the initial period, repeated 200 times. Subsequently, the query set follows a 50-way 5-shot with 50 repetitions, and the test set comprises 5 data points.

IV. EXPERIMENTAL STUDY

In this section, we evaluate the proposed approach with real-world and public IRIDIUM satellite data [10] including 66 satellites, each with 48 antennas. In our experiments, we randomly choose 10 satellites and 5 antennas from each of the satellites. Then, 50 labels in the dataset are then identified. For each label, 1,000 pieces of I/Q data are randomly selected from the original dataset. The created dataset contains 50,000 pieces of data. All the experimental environments are based on an NVIDIA A100 GPU Server and PyTorch.

A. Data Generation

First, we normalize the data and take the whole dataset as the input to generate 50,000 pieces of virtual data. Fig. 3a and Fig. 3b are the original data and generated data over a single measurement including 100 I/Q samples, respectively. We can notice that there is more noise in the generated data, and the data is not concentrated.

The initial execution of the VAE yields an FID value of 5.514. Subsequent refinement of the loss function, with α and β of each loss function equal to 1.5 and 0.25, results in a decreased FID value of 3.280. Table II illustrates different FID values corresponding to different loss function weights, in which α and β are the weights of reconstruction loss

and KL divergence loss, respectively. Table II demonstrates that in most cases, the value of FID becomes smaller as the weight of reconstruction loss becomes larger. However, with $\alpha = 0.25$ and $\beta = 1$ and 1.5, the performance of FID becomes better compared to $\alpha = 1$. We believe that this is a normal fluctuation due to smaller constraints on the original data versus the generated data, where the quality of the generated data is greatly different from the original data and cannot be used for data argumentation. Both sets of generated data, with and without the adjusted loss function, are employed for data augmentation. Maintaining a constant dataset size of 50,000, generated data is integrated with real data for model training, and the trained model is then evaluated on a real test set.

Fig. 4 illustrates the correlation between the proportion of generated data in the dataset and the accuracy of classification on real data. Notably, the VAE model with the weights α and β of each loss function is 1.5 and 0.25 consistently outperforms, achieving accuracy exceeding 95%, even in worst-case scenarios. Despite impressive performance, the generated data does not match the accuracy of real data with an accuracy of 98.05%. In an exclusive generated data scenario, the classification accuracy is 96.72%, indicating the impact of added Gaussian noise and information loss. For data augmentation, this introduced noise contributes to increased variability.

In this experiment, we first use the whole real data and then add the virtual data (i.e., generated data) to create a new dataset. Fig. 5 shows the classification accuracy after adding different amounts of generated data. It is noticed that we can achieve more than 99% accuracy when the real data is 20% of the overall data volume. We conduct further experiments, where the real data is removed. More specifically, when the training set only contains virtual data with five times the total amount of real data, and the test set includes 20% of the real data, the final accuracy can reach 99.80%, as illustrated in Fig. 5. Such high accuracy means the generation of virtual data can help the deep neural network learn intrinsic structures of the raw data. Through the utilization of numerous generated data, the model possibly learns noise effect, thereby enabling the learning of critical information and enhancing model performance. This validates the effectiveness of using this generated data to greatly enhance the model's ability for LEO satellite device fingerprinting.

B. Few Shot Learning

After confirming data augmentation's effectiveness on a single dataset, adapting the model to various data becomes a focal concern. Unlike the fixed dataset size in the prior experiment, we utilize augmented data in this process with a larger dataset of 50,000 samples. Fig. 6 shows results with different percentages of real data. To determine the superior distance metric function, prototypical network experiments consider both Euclidean distance and cosine similarity distance. Besides, the accuracy of prototypical networks on the original data is also included.

The results presented in Fig. 6 show the superior performance of the cosine similarity classifier over the Euclidean




Fig. 4: The influence of the weight of each loss function on classification accuracy in various cases.




Fig. 5: Classification accuracy after adding different amounts of generated data.




Fig. 6: The impact of distance metric functions on classification accuracy in various cases.

distance. In scenarios with all real data as input, the model employing the cosine similarity classifier achieves the accuracy of 97.54% on the original task and 77.02% on the second-period data. Conversely, the model utilizing the Euclidean distance achieves 95.78% accuracy on the original task and 70.58% on the subsequent data period in the same scenario. As the proportion of generated data increases, the model's performance on new tasks decreases, thus reaching a minimum accuracy of 60.74% when all inputs consist of generated data.

V. CONCLUSION

In this paper, we proposed a modified VAE with a tailored loss function to generate LEO satellite data and achieve data augmentation. Then, for the generated data, we used an FID metric to measure its diversity and similarity to the original data. Besides, a cosine similarity classifier was incorporated into prototypical networks for facilitating few-shot learning at LEO satellite data in different periods. In our experiments, we demonstrated the effectiveness of the proposed approach. The proposed model can achieve a 50-classification accuracy of 99.80% in a single period and obtain a significant improvement in classification accuracy, from 2% to 77% by using few-shot learning.

ACKNOWLEDGMENT

This work is also supported in part by the NSF (CNS-2321763, CNS-2319343, CNS-2317190, CNS-2107190, IIS-2306789, and IIS-2306791).

REFERENCES

- [1] Q. Li and W. Trappe, "Detecting spoofing and anomalous traffic in wireless networks via forge-resistant relationships," *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 4, pp. 793–808, 2007.
- [2] A. Jagannath, J. Jagannath, and P. S. P. V. Kumar, "A comprehensive survey on radio frequency (rf) fingerprinting: Traditional approaches, deep learning, and open challenges," *Computer Networks*, vol. 219, p. 109455, 2022.
- [3] T. Jian, B. C. Rendon, E. Ojuba, N. Soltani, Z. Wang, K. Sankhe, A. Gritsenko, J. Dy, K. Chowdhury, and S. Ioannidis, "Deep learning for RF fingerprinting: A massive experimental study," *IEEE Internet of Things Magazine*, vol. 3, no. 1, pp. 50–57, 2020.
- [4] T. Zhao, N. Wang, Y. Wu, W. Zhang, and X. Wang, "Backdoor attacks against low-earth orbit satellite fingerprinting," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2024, pp. 1–6.
- [5] T. Zhao, X. Wang, J. Zhang, and S. Mao, "Explanation-guided backdoor attacks on model-agnostic rf fingerprinting," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*. IEEE, 2024, pp. 1–10.
- [6] T. Zhao, X. Wang, and S. Mao, "Cross-domain, scalable, and interpretable rf device fingerprinting," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*. IEEE, 2024, pp. 1–10.
- [7] Z. Qu, G. Zhang, H. Cao, and J. Xie, "LEO satellite constellation for Internet of Things," *IEEE access*, vol. 5, pp. 18 391–18 401, 2017.
- [8] M. Motalebighomi, H. Sathaye, M. Singh, and A. Ranganathan, "Cryptography is not enough: Relay attacks on authenticated gnss signals," *arXiv preprint arXiv:2204.11641*, 2022.
- [9] W. Liu, S. Zheng, Z. Deng, K. Wang, W. Lin, J. Lei, Y. Jin, and H. Liu, "Multi-scene Doppler power spectrum modeling of LEO satellite channel based on atlas fingerprint method," *IEEE Access*, vol. 9, pp. 11 811–11 822, 2021.
- [10] G. Oliveri, S. Sciancalepore, and R. Di Pietro, "Physical-layer data of IRIDIUM satellites broadcast messages," *Data in Brief*, vol. 46, p. 108905, 2023.
- [11] G. Oliveri, S. Sciancalepore, S. Raponi, and R. Di Pietro, "PAST-AI: Physical-layer authentication of satellite transmitters via deep learning," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 274–289, 2022.
- [12] B. Ding, L. Liu, L. Bing, C. Kruengkrai, T. H. Nguyen, S. Joty, L. Si, and C. Miao, "DAGA: Data augmentation with a generation approach for low-resource tagging tasks," *arXiv preprint arXiv:2011.01549*, 2020.
- [13] J. Tomczak and M. Welling, "VAE with a VampPrior," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 1214–1223.
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [15] M. Patel, X. Wang, and S. Mao, "Data augmentation with conditional GAN for automatic modulation classification," in *Proceedings of the 2nd ACM Workshop on wireless security and machine learning*, 2020, pp. 31–36.
- [16] C. Tarn and W.-F. Zeng, "pDeep3: toward more accurate spectrum prediction with fast few-shot learning," *Analytical Chemistry*, vol. 93, no. 14, pp. 5815–5822, 2021.
- [17] S. Mackey, T. Zhao, X. Wang, and S. Mao, "Poster abstract: Cross-domain adaptation for RF fingerprinting using prototypical networks," in *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, 2022, pp. 812–813.
- [18] H. Mazar, *Radio spectrum Management: Policies, regulations and techniques*. John Wiley & Sons, 2016.
- [19] J. Liu, L. Song, and Y. Qin, "Prototype rectification for few-shot learning," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*. Springer, 2020, pp. 741–756.