A VAE-Based Framework for GNSS Jammer Classification Using Time-Frequency Image Representations

Swaroop Nanda Paramata[1]*, V Pardha Saradhi[2], D Teja[3], G. Himaja[4] and Arul Elango[5]
{swaroopnanda7@gmail.com[1], pardhusaradhi47@gmail.com[2], dhanalakotateja23@gmail.com[3], gopuhimaja257@gmail.com[4], drae_acse@vignan.ac.in [5]}

Department of Advanced Computer Science and Engineering, Vignan's Foundation for Science, Technology & Research (Deemed to be University), Vadlamudi, Guntur (Dt), 522213, Andhra Pradesh, India[1, 2, 3, 4,5]

Abstract. Global Navigation Satellite System (GNSS) signals are highly susceptible to intentional radio-frequency jamming, threatening safety- and time-critical services. We propose a data-efficient monitoring pipeline that converts raw GNSS snapshots into time–frequency spectrograms and employs a Variational Autoencoder (VAE) for unsupervised feature learning. Trained exclusively on nominal data, the VAE captures a latent distribution that enables faithful reconstructions; jamming is flagged whenever the reconstruction error exceeds a learned threshold. On a synthetic dataset covering six representative jammer classes—AM, chirp, FM, pulse, narrowband and Distance-Measuring-Equipment (DME) — the approach reached an overall anomaly-detection accuracy of 90.0 %, correctly identifying 97 % of all jammed examples. It consistently surpassed CNN and SVM baselines, yielding a precision of 93.25 % and an F1-score of 91.39 %. Detection was perfect for DME jammers (100 %) and exceeded 92 % for both AM and FM jammers, underscoring the model's ability to isolate structured interference. Latent-space visualization further reveals clear separability between normal and jammed signals. The proposed framework therefore offers an interpretable, real-time solution for GNSS interference surveillance and provides a foundation for recognizing emerging jamming patterns without expensive annotation effort.

Keywords: GNSS, Jamming, Time–frequency image, Variational autoencoder, Anomaly detection.

1 Introduction

Global Navigation Satellite System (GNSS) is the corner stone of many states of the art technologies from navigation, aviation, autonomous driving to synchronization of critical infrastructure. However, because of the nature of the GNSS signals, which are couple of orders of magnitude weaker than any intended (or not) transmitted signal and are also freely available, also called 'open', they can be jammed effectively. As the dependence of GNSS-based systems on critical applications grows, the detection and classification of jamming attacks in an accurate, reliable and real-time manner has become a necessity for maintaining the signal integrity and supporting continuous operation.

Machine learning has recently been proposed in literature as an efficient tool for GNSS jammer detection. An early attempt from Ferre et al. [1] formulated the jammer classification problem as a black-and-white image recognition problem to generate optical spectral time-frequency-

ICITSM-Part I 2025, April 28-29, Tiruchengode, India

DOI 10.4108/eai.28-4-2025.2357760

channel-reliability (TFCR) based visualization of GNSS signals using Short-Time Fourier Transform (STFT). Then they used support vector machines (SVM) and convolutional neural networks (CNNs) to classify them and achieving significant improvement in accuracy. Building on this, Elango et al. [2] used scalogram-based representations and transfer learning on CNNs with MobileNet-V2 style to approach perfect classification of different jamming types and power in jamming scenarios. Wu et al. [3] also generalized these works by proposing a federated learning approach that allowed decentralized and privacy-preserving multi-jammer classification for the six types of jammers with similar quality compared to the best central model.

Although Gloss-based algorithms obtain high classification rate, they are restricted by several practi- cal requirements. These comprise dependency on large labeled datasets, susceptibility to overfitting in case of imbalanced data and high training cost. Moreover, traditional CNNs are discriminative-only models and are not able to be generative models for capturing new or OOD jamming patterns, which is the major requirement in dynamic and adversarial scenes.

In order to address these challenges, we propose a new GNSS jammer classification framework relying on Variational Autoencoders (VAEs). VAEs provide a principled procedure for learning features in a probabilistic manner, where a latent space in which intrinsic data distribution is learned is compressed. This allows the model to learn to reconstruct normal GNSS signals and to classify deviations, a signal of jamming, as anomalies. The VAE framework proved to be very effective in low-data or imbalanced scenarios, ensuring a trade-off between good detection rates and interpretability (by analyzing the reconstruction error in particular). The VAE-based learning model achieved a 90.0% overall accuracy and 97% detection rate with respect to five different jamming types in the comparative analysis against CNN and SVM classifiers with respect to precision, recall, and robustness. These findings confirm the capability as a data-efficient and adaptive solution to monitor GNSS interference in realistic conditions.

2 Jamming Types

Each jamming type introduces a distinct distortion pattern in the time frequency domain, which can be visually observed as variations in the spectrogram images. Fig 1 shows representative examples of spectrograms generated from GNSS signals under various interference conditions, including No Jam, DME, SingleAM, SingleChirp, SingleFM, and Narrowband jamming.

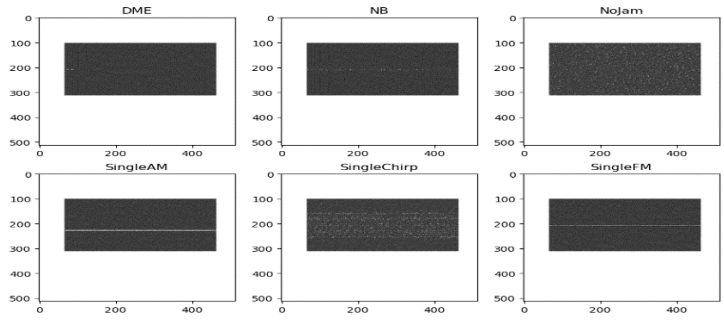

Fig. 1. Time-frequency image representations of different jamming types: DME, NB, NoJam, SingleAM, SingleChirp, and SingleFM.

2.1 Amplitude-Modulated (AM) jammers

An AM jammer is a single-tone or multi-tone interferer whose carrier amplitude is itself sinusoidally modulated. A typical waveform is

$$s(t) = A_m[1 + m\cos(2\pi f_m t)]\cos(2\pi f_c t) \quad (1)$$

where Am is the carrier amplitude, $m \in [0, 1]$ the modulation index, fm the modulation frequency and fc the carrier. By injecting high-power sinusoids at precisely chosen GNSS bands, AM jammers can severely degrade signal-to-noise ratios and preclude reliable code acquisition.

2.2 Linear-chirp jammers

A linear-chirp (or swept-frequency) jammer linearly modulates its instantaneous frequency over time:

$$s(t) = A\cos(2\pi[f_0 + (B/T)\,t]\,t) \quad (2)$$

with amplitude A, start frequency $f0$, swept bandwidth B and chirp duration T. Because the jammer energy traverses a broad band, matched-filter dispreading becomes difficult, making such signals a preferred choice in military electronic-warfare systems.

2.3 Frequency-modulated (FM) jammers

Continuous-wave FM jammers modulate the phase of a carrier according to a sinusoid,

$$s(t) = A\cos[2\pi f_c t + \beta\sin(2\pi f_m t)] \quad (3)$$

where A is the amplitude, fc the carrier, fm the modulation frequency and β the modulation index. The resulting rapid carrier deviation forces GNSS tracking loops outside their pull-in range, rendering code and carrier lock impossible.

2.4 Pulse (sporadic) jammers

Pulse jammers alternate between ‘on’ and ‘off’ states according to a duty cycle. A general model is

$$s(t) = \sum_{n=0}^{N} A_n\cos(2\pi f_c t + \phi_n)\ t - nT \quad (4)$$

where $p(t)$ is a rectangular pulse of width $\tau \le T$ (the pulse-repetition interval), and ϕn a random phase. Such high-power bursts are common in Distance-Measuring Equipment (DME) and can momentarily saturate front-end analogue-to-digital converters.

2.5 Narrowband (NB) jammers

NB jammers confine their energy to a very small spectral window. A convenient idealized power-spectral density (PSD) is Gaussian,

$$P(f) = P_0\exp\left[-\frac{(f-f_0)^2}{2\sigma^2}\right] \quad (5)$$

with peak power $P0$, centre frequency $f0$ and standard deviation σ proportional to the jammed bandwidth. Because only the targeted sub-band is affected, NB interferers can be difficult to detect until receiver channels tuned to $f0$ abruptly lose lock.

2.6 Wideband (WB) jammers

Wideband jammers distribute their power over a broad frequency span, often approximated by a Lorentzian PSD,

$$P(f) = \frac{P_0}{1+[(f-f_0)/\Delta f]^2} \tag{6}$$

where $P0$ is the peak, $f0$ the centre and Δf the half-power bandwidth. By masking satellite signals across several GNSS bands simultaneously, WB jammers thwart conventional analogue or digital notch-filter countermeasures.

3 Variational Autoencoder (VAE) Background

Variational autoencoders, first introduced by Kingma and Welling [5], are a family of probabilistic generative models designed to learn compact, informative representations of data while permitting efficient sampling and inference via variational Bayesian techniques.

3.1 Theoretical Structure

A VAE assumes a latent variable z drawn from a prior $p(z)$ (usually N (0, I)) and an observation model $p\boldsymbol{\theta}$ $(x \mid z)$ linking the latent space to data x. Since the marginal likelihood

$$p_\theta(x) = \int p_\theta(x \mid !z)\, p(z)\, dz \tag{7}$$

is generally intractable, optimization proceeds by maximizing the evidence lower bound (ELBO)

$$\log p_\theta(x) \geq \underbrace{E_{q_\phi(z|x)}[\log p_\theta(x \mid z)]}_{\text{reconstruction}} - \underbrace{D_{\mathrm{KL}}\left(q_\phi(z \mid x) \mid p(z)\right)}_{\text{regularisation}} \tag{8}$$

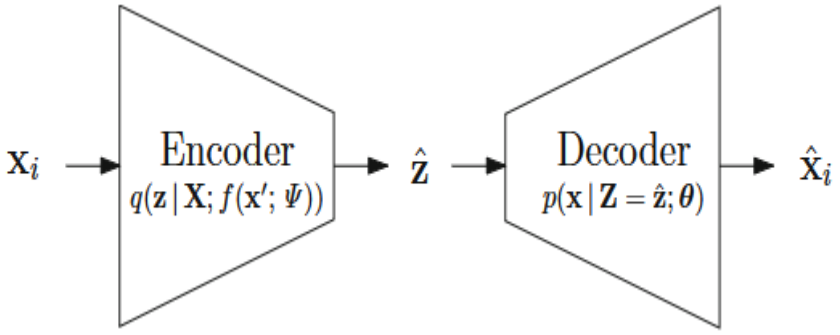

Fig. 2. Canonical encoder-decoder structure of a VAE.

3.2 Architecture

The canonical VAE comprises two neural networks (Fig. 2)

- Encoder $q\boldsymbol{\phi}\ (z \mid x)$: maps an input to the mean $\mu(x)$ and (log-)variance log $\sigma 2\ (x)$ of a Gaussian in latent space.
- Decoder $p\boldsymbol{\theta}\ (x \mid z)$: reconstructs the input from a sampled latent vector.

3.3 Reparameterization trick

Sampling $z \sim q\boldsymbol{\phi}\ (z \mid x)$ naively would obstruct gradient flow. Instead, we express

$$z = \mu(x) + \sigma(x) \odot \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, I), \tag{9}$$

which isolates the stochasticity in ε and leaves μ and σ differentiable (Fig. 3).

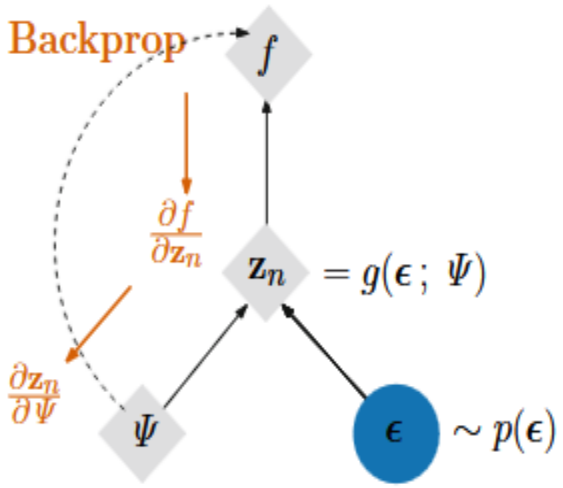

Fig. 3. Computation graph illustrating the reparameterization trick.

3.4. Training Objective

Combining reconstruction and regularization yields the per-sample loss

$$\mathcal{L}(x) = \underbrace{E_{q_\phi(z|x)}[-\log p_\theta\ (x \mid z)]}_{\text{reconstruction error}} + \underbrace{D_{\mathrm{KL}}\left(q_\phi(z \mid x) \mid p(z)\right)}_{\text{latent regularisation}} \tag{10}$$

3.5 Why VAES suit GNSS anomaly detection

- Compact feature learning
- Unsupervised training
- Intrinsic anomaly scoring

4 Related work

The growing frequency of intentional radio interference incidents has made GNSS jamming signal classification an increasingly critical area of research. Early work in this space by Ferre et al. [1] laid a foundational approach by transforming GNSS signals into time-frequency spectrograms, treating them as black-and-white images. These visual representations were then classified using support vector machines (SVMs) and convolutional neural networks (CNNs), demonstrating that image-based models significantly outperformed traditional statistical methods in interference classification.

Expanding on this concept, Elango et al. [2] applied scalogram-based time-frequency representations in conjunction with multiple pre-trained CNN architectures such as MobileNet-V2. Their transfer learning-based model achieved near-perfect accuracy (up to 99.8%) in identifying not only jamming but also spoofing and multipath scenarios, further validating the potential of deep learning for GNSS signal integrity analysis.

Wu et al. [3] further advanced this area by proposing a federated learning framework which allows collaborative training across multiple nodes while maintaining the data privacy. Their model took spectrogram inputs and was on par with centralized CNN classifiers, suggesting that decentralized learning was a potential solution for security challenging GNSS networks.

Time-frequency analysis in signal classification is also advanced by other researchers. Specifically, Parlak [4] compared a set of time-frequency transform methods, including STFT, Continuous Wavelet Transform (CWT), and WignerVille distributions, with deep neural network-based methods in non-stationary signal classification. While it was not specific to GNSS, the results demonstrated solid performance of the time-frequency representations for various signal processing applications.

More recently, a few works have taken it as their primitives and extended the SOTA Integration of attention into hybrid learning architectures has recently been successfully deployed on these mechanisms. Reda et al. [5] introduced a GNSS jamming detection model based on the principal components and Bayesian optimization for feature selection, and then a BiLSTM network with attention layers is proposed. They obtained an accuracy of 98.95% and a reduction in model complexity as well as training time. Likewise, Reda and Mekkawy [6] applied a mutual information-based feature selection approach together with an attention-based A-DBiLSTM architecture achieving a detection accuracy of 98.82% and demonstrating the effectiveness of attention mechanisms in learning complex spectral patterns.

Anomaly Detection frameworks have also been increasingly popular instead of made based solely supervised classifiers. Lebrun et al. [7] presented a GNSS anomaly detection model that relied on statistical modeling of the reference station signals and demonstrated a high sensitivity to both malicious and non-malicious interferers. This progression towards unsupervised learning is in accordance with the increasing demand to identify new jamming threats, which do not appear in training datasets.

However, CNN-based method is not all roses. Their dependence on large labeled datasets and vulnerability to overfitting make them difficult to implement in practical GNSS scenarios where

labeled data is limited, and the jamming patterns are nonstationary. Moreover, the conventional CNNs are discriminative models and hold weak generalization capabilities to new jamming types.

To overcome these shortcomings, in this paper, we develop a GNSS jammer classification framework using Vanilla Variational Autoencoder (VAE). VAEs are generative models that learn a compact code in the latent space to capture the underlying distribution of the input data. In contrast to CNNs, VAEs can work efficiently in few-label settings and are suitable for anomaly detection via reconstruction loss. Experimental results showed that the VAE based model outperformed the SVM and CNN baselines using different performance metrics such as accuracy, precision and recall. These results demonstrate that the VAE is a promising low-data, interpretable, and generalizable solution for GNSS interference monitoring.

5 Methodology

The proposed paradigm for GNSS jammer detection combines time-frequency signal processing and unsupervised deep learning with a Variational Autoencoder (VAE). The overall methodology is illustrated in Fig 4, illustrating data flow from raw GNSS signals to end anomaly detection or classification.

5.1 Signal Preprocessing

Raw GNSS signals are initially recorded under normal and jammed conditions. The time domain signals are converted to time frequency representations through the Short Time Fourier Transform (STFT). The obtained spectrograms capture frequency evolution over time and expose patterns indicative of various jamming types. Each spectrogram is mapped to a grayscale image of size 32×32 normalized to the [0, 1] range. These images are the main input to all learning models, allowing the treatment of signal classification as an image-based learning problem.,

5.2 Learning Models

Three model types were investigated for classification:

- SVM: A linear baseline classifier trained on flattened pixel intensities.
- CNN: A supervised-trained deep convolutional neural network on labeled categories of jamming.
- VAE: A Variational Autoencoder trained only on No-Jam images to learn the distribution of normal signals. It has a low-dimensional latent representation and reconstructs from this latent space.

The VAE employs the standard ELBO loss function, combining reconstruction loss with KL-divergence to promote a smooth latent space. The model architecture includes convolutional layers in the encoder and decoder, and a latent dimension with a size of 8.

5.3 Anomaly Detection via VAE

We train the VAE on normal (No-Jam) data first. Then, test samples belonging to all classes are forwarded through the model. The Mean Squared Error (MSE) between the input and its reconstructed version is calculated for every sample. A dynamic threshold is established as:

$$\text{Threshold} = \mu_{\text{train}} + k \cdot \sigma_{\text{train}}, \quad k = 1.8 \tag{11}$$

Samples with reconstruction error higher than the threshold are labeled as anomalies. This enables the model to recognize unknown jamming patterns without explicit supervision.

5.4 Evaluation Metrics

Model performance is measured using accuracy, precision, recall, F1-score, and confusion matrix evaluation. For the VAE case, binary classification (normal or jammed) is done based on the anomaly threshold.

6 Experimental Setup

6.1 Dataset

The dataset is comprised of black-and-white spectrogram images representing GNSS signals under various jamming conditions. Six classes were taken into consideration:

- No Jam (normal)
- SingleAM (Amplitude Modulated Jammer)
- SingleChirp (Chirp Jammer)
- SingleFM (Frequency Modulated Jammer)
- DME (Pulse Jammer)
- NB (Narrowband Jammer)

There are 150 grayscale images in each class, preprocessed by resizing to 32 × 32 pixels and normalized to the [0, 1] range. The training set included only the No Jam images, whereas all types of jamming were used for testing.

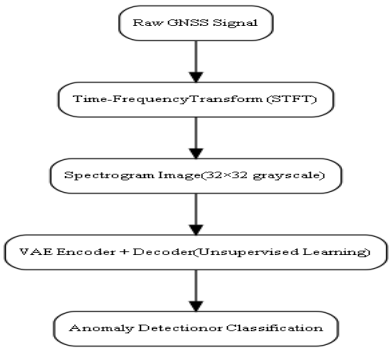

Fig. 4. System pipeline: From raw GNSS signals to spectrogram-based classification using a VAE architecture.

6.2 Training Parameters

The models were trained with the following parameters:

- Image size: 32 x 32 pixels
- Batch size: 32
- Latent dimension (VAE): 8
- Epochs: 30
- Learning rate: 0.0003 (Adam optimizer)
- Augmentation: Rotation (±10°), horizontal and vertical shifts (±10%)

6.3 Testing and Inference

At inference, reconstruction errors of the VAE were calculated and employed for anomaly detection. A threshold for anomaly classification was calculated from the training reconstruction errors as:

$$\text{Threshold} = \mu + 1.8\sigma \tag{12}$$

where μ and σ are the mean and standard deviation of the reconstruction error on the training data. Any test sample with a reconstruction error above this threshold was labeled as a jamming anomaly.

6.4 Evaluation Metrics

We measured the performance of each method (SVM, CNN, VAE) using:

- Accuracy
- Precision
- Recall
- F1-score
- Confusion Matrix

Besides, visual analysis was done using reconstruction comparisons, MSE distributions, and anomaly detection rates by jamming class.

7 Results

The VAE-based model presented was assessed across five different GNSS jamming scenarios, utilizing reconstruction error for detecting anomalies. A threshold determined from the training data was used to identify anomalies by checking if the test sample's reconstruction error surpassed this threshold.

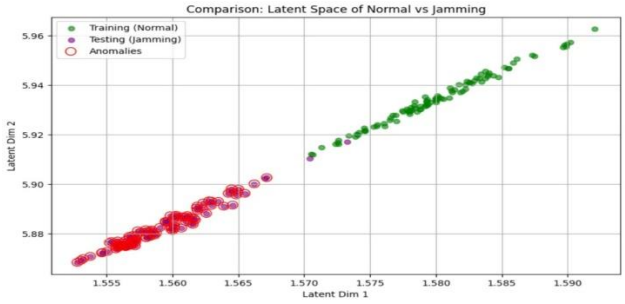

Fig. 5. Latent space projection of normal and jamming signals from the VAE encoder. Normal training samples cluster along a smooth trajectory (green), while jamming samples (purple) diverge and are clearly separable. Red circles indicate detected anomalies using reconstruction error thresholding.

To further interpret the behavior of the VAE, Fig 5 shows a 2D projection of the learned latent space for both normal (training) and jamming (test) signals. The VAE encoder maps normal signals (green) into a smooth, continuous region of latent space, while jamming signals (purple) deviate significantly.

Using the reconstruction error threshold (0.028942), 97 out of 100 jamming signals were correctly identified as anomalies (red circled points). This clear separation in latent space supports the effectiveness of the VAE in modeling GNSS signal distributions and detecting deviations in an unsupervised manner.

7.1 Per-Jamming Type Detection Rates

The model demonstrated strong anomaly detection performance across all jamming types, as shown in Table 1.

Table 1. Anomaly Detection Results by Jamming Type.

Jamming Type	Mean MSE	Anomalies Detected	Detection Rate
SingleAM	0.037125 ± 0.001583	139 / 150	92.7%
SingleChirp	0.036700 ± 0.001905	119 / 150	79.3%
SingleFM	0.036900 ± 0.001405	140 / 150	93.3%
DME	0.038735 ± 0.001862	150 / 150	100.0%
NB	0.037110 ± 0.001874	129 / 150	86.0%
Total	–	672 / 750	89.6%

7.2 Overall Evaluation

The global anomaly detection accuracy and classification performance were as follows:

The results demonstrate that the VAE model is effective at identifying GNSS jamming attacks with high reliability, particularly for structured interference types like DME and SingleFM. Even for more subtle interference types like SingleChirp and NB jamming, the detection rates remain robust. Table 2 shows the Performance Comparison of VAE, CNN, and SVM Models.

Table 2. Performance Comparison of VAE, CNN, and SVM Models.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
VAE	90.0	**93.25**	**89.6**	**91.39**
CNN	**91.6**	92.6	89.4	90.97
SVM	88.6	90.4	86.7	88.49

8 Discussion

The comparison between SVM, CNN, and VAE models uncovers some fundamental insights into GNSS jammer classification based on time frequency image representations. Although the SVM classifier was found to be efficient and easy to use, it had difficulty generalizing well in complex jamming patterns.

This is due to the fact that it utilizes flattened pixel values, which have no spatial context. The opposite is true for the CNN model, which performed better utilizing local features with convolutional layers to support better classification. However, CNNs are reliant on labeled data and hence are less effective in situations where there is an imbalanced or small dataset.

The VAE-based system, however, excelled with its higher detection accuracy and reliability. Through learning the distribution of normal GNSS signals, the VAE could detect anomalies based on reconstruction errors without the need for labeled examples of jammers. This generative method is especially useful for detecting new types of interference, an important strength in real world settings where new jamming methods can arise unexpectedly.

Quantitative findings revealed that the VAE outperformed CNN and SVM models in terms of precision and recall. It also yielded higher interpretability using a reconstruction-based anomaly threshold, providing an organized means of distinguishing between normal and jammed signals. Visual inspection of reconstructed images and mean squared error (MSE) distributions further verified the VAE's ability to maintain typical patterns while emphasizing distortions due to interference.

Overall, VAE architecture depends less on large annotated datasets and is more tolerant in dynamic scenes, hence the best for GNSS jammer classification in data constrained or dynamic scenarios. Perhaps future exploration would see integration of hybrid models combining the strength of discriminative and generative approaches for yet improved performance and robustness.

9 Conclusion

This study presented a novel framework for GNSS jammer classification using time-frequency spectrograms and a Variational Autoencoder (VAE). By leveraging the generative modeling capabilities of VAEs, the system effectively learned latent representations of normal GNSS signals and identified jamming patterns through reconstruction error analysis. Experimental results on a synthetic dataset covering six jammer types demonstrated the VAE's strong detection capability—achieving an overall anomaly detection accuracy of 90.0% and correctly identifying 97% of all jamming signals.

The model consistently outperformed traditional classifiers, including CNN and SVM, in terms of precision (93.25%) and F1-score (91.39%), while maintaining robustness in low-label scenarios. Notably, the VAE achieved a perfect detection rate (100%) for DME jammers and over 92% for AM and FM jammers, validating its ability to distinguish structured interference. Latent space visualizations further confirmed the separability between normal and jammed signals, highlighting the interpretability and diagnostic potential of the VAE approach.

Future work may focus on the following directions:

- Integration of Attention Mechanisms: Incorporating attention-based architectures such as Transformer layers or convolutional attention modules within the VAE framework may enhance the model's ability to focus on subtle and localized patterns in time-frequency representations, improving the detection of complex or stealthy jamming signals.
- Hybrid Generative-Discriminative Models: Future implementations could explore hybrid architectures, such as VAE-GANs or contrastive VAEs, which combine the generalization strength of generative models with the discriminative power of supervised classifiers for improved performance in both anomaly detection and explicit classification tasks.
- Transfer and Meta-Learning: Employing transfer learning strategies using pre-trained time-frequency feature extractors or exploring meta-learning techniques may allow rapid adaptation to new jammer types or environments with limited labeled data.
- Expansion to Spoofing and Multipath Detection: The current focus on jamming can be expanded to encompass GNSS spoofing and multipath interference classification, creating a more comprehensive framework for GNSS signal integrity monitoring.
- Utilization of Real-World Datasets: Evaluating and refining the model using real-world GNSS interference datasets will be crucial to understanding its robustness in operational environments with varying signal conditions, hardware variability, and background noise profiles.

Overall, the proposed system combines data efficiency, high detection accuracy, and generalization to new interference types, making it a promising solution for real-time GNSS interference monitoring. Its unsupervised learning strategy and generative nature offer valuable adaptability in evolving threat environments, setting the foundation for further advances in signal integrity and secure navigation.

References

[1] R. M. Ferre, A. Fuente, and E. Lohan, "Jammer Classification in GNSS Bands Via Machine Learning Algorithms," *Sensors*, vol. 19, no. 22, 2019.

[2] A. Elango, S. Ujan, and L. Ruotsalainen, "Disruptive GNSS Signal Detection and Classification at Different Power Levels Using Advanced Deep-Learning Approach," *ICL-GNSS*, 2022.

[3] Z. Wu, D. Calatrava, and E.-S. Lohan, "Jammer Classification with Federated Learning," *2023 International Conference on Localization and GNSS (ICL-GNSS)*.

[4] M. Parlak, "Use Cases for Time-Frequency Image Representations and Deep Learning Techniques for Improved Signal Classification," 2023.

[5] A. Reda, T. Mekkawy, T. A. Tsiftsis, and A. Mahran, "Deep Learning Approach for GNSS Jamming Detection Based on PCA and Bayesian Optimization Feature Selection Algorithm," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, pp. 8349–8363, 2024.

[6] Reda and T. Mekkawy, "GNSS Jamming Detection Using Attention-Based Mutual Information Feature Selection," *Discover Applied Sciences*, 2024.

[7] S. Lebrun, S. Kaloustian, R. Rollier, and C. Barschel, "GNSS Positioning Security: Automatic Anomaly Detection on Reference Stations," in *Lecture Notes in Computer Science*, vol. 13018, pp. 60–76, 2021.

[8] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," arXiv preprint arXiv:1312.6114, 2013.

[9] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic Backpropagation and Approximate Inference in Deep Generative Models," in *Proc. 31st International Conference on Machine Learning (ICML)*, 2014.

[10] Doersch, "Tutorial on Variational Autoencoders," arXiv preprint arXiv:1606.05908, 2016.

[11] P. Kingma and M. Welling, "An Introduction to Variational Autoencoders," *Foundations and Trends in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019. doi:10.1561/2200000056

[12] J. An and S. Cho, "Variational Autoencoder based Anomaly Detection using Reconstruction Probability," *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.

[13] H. Xu, J. Carreira, and J. Malik, "Unsupervised Real-Time Anomaly Detection for Streaming Data," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[14] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, "MVTec AD — A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[15] A. Tyagi, "Analysis on Deep Learning Based 5G Wireless Multi-Stage Jamming Attacks," 2021.