

LATENT-SPACE METRICS FOR COMPLEX-VALUED VAE OUT-OF-DISTRIBUTION DETECTION UNDER RADAR CLUTTER

Y. A. Rouzoumka^{1,2} E. Terreaux¹ C. Morisseau¹ J.-P. Ovarlez^{1,2} C. Ren²

¹ DEMR, ONERA, Université Paris-Saclay, 91120 Palaiseau, France

² SONDRRA, CentraleSupélec, Université Paris-Saclay, 91190 Gif-sur-Yvette, France.

ABSTRACT

We investigate complex-valued Variational AutoEncoders (CVAE) for radar Out-Of-Distribution (OOD) detection in complex radar environments. We proposed several detection metrics: the reconstruction error of CVAE (CVAE-MSE), the latent-based scores (Mahalanobis, Kullback-Leibler divergence (KLD)), and compared their performance against the classical ANMF-Tyler detector (ANMF-FP). The performance of all these detectors is analyzed on synthetic and experimental radar data, showing the advantages and the weaknesses of each detector.

Index Terms— Complex-Valued VAE, Out-Of-Distribution detection, Latent Metrics Detection, Radar Detection.

1. INTRODUCTION

Classical adaptive radar detectors (Kelly’s detector, AMF, ANMF) rely on Gaussian and SIRV clutter models and may degrade under non-Gaussian plus additive thermal noise backgrounds [1, 2, 3]. We therefore pursue a data-driven alternative based on complex-valued VAEs and latent-space OOD scores.

In recent years, data-driven approaches have emerged to alleviate the need for precise clutter modeling. Among them, VAEs [4] have demonstrated promising capabilities for anomaly and OOD detection in diverse applications, including radar detection [5], speech enhancement [6], medical imaging [7], industrial monitoring [8], and acoustic signal analysis [9]. These models learn a latent representation of the training data and use reconstruction or probabilistic criteria to detect deviations. Despite their effectiveness, most VAE-based detectors operate in the real domain and often treat complex-valued radar data by separating real and imaginary components into distinct channels.

Recent advances in Complex-Valued Neural Networks (CVNNs) have shown the benefits of directly modeling complex-valued signals [10, 11]. In radar signal processing, complex-valued models are particularly well-suited due to the

inherent phase information contained in the data. Complex-Valued VAEs (CVAEs) extend this paradigm, enabling phase-aware latent modeling and better structural encoding of radar signals.

However, while CVAEs offer promising avenues for radar detection, the literature lacks a comprehensive evaluation of how different OOD detection strategies behave in the complex latent space. Most existing works rely on simple reconstruction error metrics (e.g., MSE) in the observation domain, without exploiting the statistical richness of the learned latent distribution. Moreover, it remains unclear whether latent-based scores such as Mahalanobis distance [12], KLD, or magnitude-based norms provide robust alternatives for radar detection, especially under challenging noise conditions (e.g., compound clutter or low SNR regimes).

We present the study of complex-valued latent-space OOD scores learned by a CVAE for radar, and position them against a classical detector. Our contributions are (i) benchmark of CVAE latent-space OOD scores (Mahalanobis and empirical KLD) versus the reconstruction error (CVAE-MSE) and a strong classical baseline (ANMF-Tyler built with Tyler estimate) [13], all calibrated at $P_{fa} = 10^{-2}$, (ii) evaluate across cGN+AWGN, cCGN+AWGN, and experimental CSIR data, reporting both the mean over cells (excluding Doppler 0) and the specific Doppler-0 cell, (iii) analyze robustness vs. SNR: (1) CVAE-MSE leads on synthetic data; (2) ANMF-FP is particularly effective under compound-Gaussian clutter and on real data; (3) Hermitian Mahalanobis mitigates Doppler-0 degradation relative to CVAE-MSE, (iv) KLD is brittle in practice and (iv) give guidelines on selecting reconstruction, latent, or classical scores depending on clutter realism and cell-specific behavior (e.g., Doppler 0).

Notations: Matrices are in bold and capital, vectors in bold. For any matrix $\mathbf{A}$ or vector, $\mathbf{A}^T$ is the transpose of $\mathbf{A}$, $\mathbf{A}^*$ the (elementwise) complex conjugate, and $\mathbf{A}^H = (\mathbf{A}^*)^T$ the Hermitian transpose of $\mathbf{A}$. $\mathbf{I}$ is the identity matrix. $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Gamma})$ and $\mathcal{CN}(\boldsymbol{\mu}, \boldsymbol{\Gamma})$ are respectively real and complex circular Normal distribution of mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Gamma}$. The matrix operator $\mathcal{T}(\cdot)$ is the Toeplitz matrix operator $\rho \rightarrow \{\mathcal{T}(\rho)\}_{i,j} = \rho^{|i-j|}$. The symbols $\odot$ and $\oslash$ denote the Hadamard element-wise product and division, respectively. On vectors, the operator $\circ$ acts on each of its com-

Part of this work was supported by ANR-ASTRID NEPTUNE 3 (ANR-23-ASM2-0009).

ponents (power, log, modulus).

2. RELATED WORK

VAEs are widely used for unsupervised OOD detection. A common baseline flags samples with large reconstruction error (or reconstruction probability) as anomalies [14]. Beyond reconstruction, latent-space scores are often employed: the Kullback-Leibler divergence between the posterior $q(z|x)$ and the prior $p(z)$ [15], and the Mahalanobis distance computed from deep features or latent embeddings [16]. These methods remain foundational in recent comparisons of VAE-based anomaly/OOD detectors [17].

Complex-valued deep learning natively matches radar data. Foundational work on complex-valued networks established the building blocks and benefits for structured signals [10]. Despite this momentum, complex-valued VAEs applied to radar remain scarce. To the best of our knowledge, prior work in adjacent signal-processing domains (e.g., speech enhancement) demonstrates the practicality of complex-valued VAEs, but radar-specific CVAE detectors are still under-explored; our study fills this gap by benchmarking latent-space criteria in complex-valued settings.

Latent-space scoring has a long pedigree: the Mahalanobis distance [18] adapted to deep features offers a class-conditional statistical deviation measure [16], while the VAE’s KLD term provides a natural posterior-prior mismatch signal [15]. However, most of these investigations focus on *real-valued* latents; their behavior in complex-valued generative models (and in radar specifically) remains comparatively underexplored, motivating the contributions of this work.

3. METHODOLOGY

This section presents our comprehensive framework for evaluating latent space detection methods in complex-valued VAEs. We first describe the complex-valued VAE architecture and its adaptations for radar signal processing [13], then provide detailed mathematical formulations and methodological developments for our two primary contributions: KL divergence-based detection and Hermitian Mahalanobis distance in complex latent spaces.

3.1. Complex-Valued VAE Architecture

Our approach builds upon the complex-valued VAE architecture specifically designed for radar signal processing applications. Unlike conventional real-valued VAEs that require separate processing of real and imaginary components, our model operates directly in the complex domain, preserving crucial phase information inherent in radar signals.

The encoder maps high-dimensional complex radar signals $\mathbf{x} \in \mathbb{C}^N$ into a lower-dimensional latent representation through complex-valued convolutional layers with complex

batch normalization and CReLU activation functions. The encoder outputs three key parameters: the mean vector $\boldsymbol{\mu} \in \mathbb{C}^q$, the variance vector $\boldsymbol{\sigma}^{\circ 2} \in \mathbb{R}^{q+}$, and the pseudo-variance term $\boldsymbol{\delta} \in \mathbb{C}^q$, which allows greater flexibility in modeling the latent space.




Fig. 1: Complex-Valued VAE network architecture

To enable efficient sampling in the complex latent space while maintaining differentiability, we introduce a novel reparameterization trick. Given the estimated parameters $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$, and $\boldsymbol{\delta}$, the latent variable $\mathbf{z}$ is sampled as follows [19]:

$$\mathbf{z} = \boldsymbol{\mu} + \mathbf{k}_r \odot \boldsymbol{\epsilon}_r + i \mathbf{k}_i \odot \boldsymbol{\epsilon}_i, \quad (1)$$

$$\text{where } \begin{cases} \mathbf{k}_r = \frac{1}{\sqrt{2}} (\boldsymbol{\sigma} + \boldsymbol{\delta}) \odot (\boldsymbol{\sigma} + \text{Re}(\boldsymbol{\delta}))^{\circ \frac{1}{2}}, \\ \mathbf{k}_i = \frac{1}{2} \left(\boldsymbol{\sigma}^{\circ 2} - (|\boldsymbol{\delta}|^{\circ})^{\circ 2} \right)^{\circ \frac{1}{2}} \odot (\boldsymbol{\sigma} + \text{Re}(\boldsymbol{\delta}))^{\circ \frac{1}{2}}. \end{cases}$$

Here, $\boldsymbol{\epsilon}_r$ and $\boldsymbol{\epsilon}_i$ are identically and independently distributed according to standard Gaussian noise vectors $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

The decoder reconstructs the radar signal $\mathbf{x} \in \mathbb{C}^N$ from the latent variable $\mathbf{z}$ using transposed complex convolutions, complex batch normalization, and CReLU activations. It mirrors the encoder in reverse to upsample and restore the original signal structure.

The model is trained by minimizing the complex-valued ELBO:

$$\mathcal{L}_{\text{CVAE}} = \mathcal{L}_{\text{rec}} + \beta \mathcal{D}_{\text{KL}}, \quad (2)$$

where the reconstruction loss is $\mathcal{L}_{\text{rec}} = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$ and the KL divergence for complex latent spaces against a standard circular complex Gaussian prior is:

$$\mathcal{D}_{\text{KL}} = \|\boldsymbol{\mu}\|^2 + \mathbf{1}_q^T \left(\boldsymbol{\sigma} - \frac{1}{2} \log^{\circ} \left(\boldsymbol{\sigma}^{\circ 2} - (|\boldsymbol{\delta}|^{\circ})^{\circ 2} \right) \right). \quad (3)$$

3.2. Detection Scores Considered

We compare three scores computed with a complex-valued VAE trained on clutter: (i) a reconstruction-based baseline (CVAE_MSE), and two latent-space scores: (ii) an empirical complex KL divergence against the marginal H_0 , and (iii) a Hermitian Mahalanobis distance in the complex latent space.

3.2.1. Reconstruction-Based Detection (CVAE_MSE)

Our baseline follows the standard CVAE paradigm: targets are identified through reconstruction error defined as

$S_{\text{MSE}}(\mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$ with $\mathbf{x}, \hat{\mathbf{x}} \in \mathbb{C}^N$. The binary decision uses a threshold λ_{MSE} calibrated on clutter (H_0) to meet a prescribed PFA:

$$S_{\text{MSE}}(\mathbf{x}) \underset{H_0}{\overset{H_1}{\gtrless}} \lambda_{\text{MSE}}. \quad (4)$$

3.2.2. KLD for Empirical H_0 complex Gaussian distribution

Rather than enforcing a theoretical prior $\mathcal{CN}(\mathbf{0}, \mathbf{I})$, we estimate in latent space an empirical complex Gaussian for clutter: $\mathcal{CN}(\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\Sigma}}_0, \hat{\boldsymbol{\Delta}}_0)$, where $\hat{\boldsymbol{\Sigma}}_0$ captures second-order covariance and $\hat{\boldsymbol{\Delta}}_0$ (diagonal) models residual non-circularity. Deviations of $q_\phi(\mathbf{z}|\mathbf{x})$ from this law are used for detection.

For clutter samples $\mathbf{x}_i$ the encoder outputs $(\boldsymbol{\mu}_i, \boldsymbol{\sigma}_i, \boldsymbol{\delta}_i)$. By setting $k_{i,\ell} = \max(\sigma_{i,\ell}^2 - |\delta_{i,\ell}|^2, \varepsilon)$ (stability clip), we can define the per-sample *circular* covariance $\boldsymbol{\Sigma}_{\text{enc}}(\mathbf{x}_i) = \text{diag}(k_{i,1}, \dots, k_{i,q})$ and :

$$\hat{\boldsymbol{\mu}}_0 = \frac{1}{N_c} \sum_{i=1}^{N_c} \boldsymbol{\mu}_i, \quad \hat{\boldsymbol{\Sigma}}_0 = \frac{1}{N_c} \sum_{i=1}^{N_c} \boldsymbol{\Sigma}_{\text{enc}}(\mathbf{x}_i), \quad (5)$$

Algorithm 1 Empirical KL Detection

Require: Encoder Enc , clutter set $\mathcal{D}_0$, test sample $\mathbf{x}$, target PFA α

- 1: For each $\mathbf{x}_i \in \mathcal{D}_0$, get $(\boldsymbol{\mu}_i, \boldsymbol{\sigma}_i, \boldsymbol{\delta}_i) \leftarrow \text{Enc}(\mathbf{x}_i)$; set $\boldsymbol{\Sigma}_1(\mathbf{x}_i) = \text{diag}(\mathbf{k}_i)$ with $k_{i,\ell} = \max(\sigma_{i,\ell}^2 - |\delta_{i,\ell}|^2, \varepsilon)$ and $\boldsymbol{\Delta}_1(\mathbf{x}_i) = \text{diag}(\boldsymbol{\delta}_i)$
 - 2: Estimate $(\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\Sigma}}_0)$ via (5) and $\hat{\boldsymbol{\Delta}}_0 = \text{diag}(\frac{1}{N_c} \sum_i \boldsymbol{\delta}_i)$
 - 3: Calibrate $\lambda_{\text{KL}} = \text{Percentile}(\{S_{\text{KL}}(\mathbf{x}_i)\}, 1 - \alpha)$ on $\mathcal{D}_0$
 - 4: On $\mathbf{x}$: compute $(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\delta}) \leftarrow \text{Enc}(\mathbf{x})$, form $\boldsymbol{\Sigma}_{\text{enc}}(\mathbf{x})$ and $\boldsymbol{\Delta}_{\text{enc}}(\mathbf{x})$, then $S_{\text{KL}}(\mathbf{x})$ via (7)
 - 5: Decide H_1 iff $S_{\text{KL}}(\mathbf{x}) > \lambda_{\text{KL}}$
-

We also estimate a diagonal pseudo-covariance as the average encoder output: $\hat{\boldsymbol{\Delta}}_0 = \frac{1}{N_c} \sum_i \text{diag}(\boldsymbol{\delta}_i)$.

For any test posterior density $\mathcal{CN}(\boldsymbol{\mu}_{\text{enc}}, \boldsymbol{\Sigma}_{\text{enc}}, \boldsymbol{\Delta}_{\text{enc}})$, the KL distance to the empirical null $\mathcal{CN}(\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\Sigma}}_0, \hat{\boldsymbol{\Delta}}_0)$ is

$$D_{\text{KL}} = \frac{1}{2} \left[\log \frac{\det \mathbf{K}_0}{\det \mathbf{K}_{\text{enc}}} + \text{tr}(\mathbf{K}_0^{-1} \mathbf{K}_{\text{enc}}) + (\mathbf{m}_0 - \mathbf{m}_{\text{enc}})^H \mathbf{K}_0^{-1} (\mathbf{m}_0 - \mathbf{m}_{\text{enc}}) - 2q \right]. \quad (6)$$

where the augmented mean and covariance are

$$\mathbf{m}_i = \begin{bmatrix} \boldsymbol{\mu}_i \\ \boldsymbol{\mu}_i^* \end{bmatrix}, \quad \mathbf{K}_i = \begin{bmatrix} \boldsymbol{\Sigma}_i & \boldsymbol{\Delta}_i \\ \boldsymbol{\Delta}_i^* & \boldsymbol{\Sigma}_i^* \end{bmatrix}, \quad i \in \{\text{enc}, 0\}.$$

The detection score is then $S_{\text{KL}}(\mathbf{x}) \underset{H_0}{\overset{H_1}{\gtrless}} \lambda_{\text{KL}}$ where

$$S_{\text{KL}}(\mathbf{x}) \equiv D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \parallel \mathcal{CN}(\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\Sigma}}_0, \hat{\boldsymbol{\Delta}}_0)), \quad (7)$$

3.2.3. Mahalanobis Distance in Complex Latent Space

For $\mathbf{z} \in \mathbb{C}^q$ with mean $\boldsymbol{\mu}_{\text{ref}}$ and Hermitian covariance $\boldsymbol{\Sigma}_{\text{ref}}$, we define the Mahalanobis distance as:

$$S_{\text{Maha}}(\mathbf{z}) = (\mathbf{z} - \boldsymbol{\mu}_{\text{ref}})^H \boldsymbol{\Sigma}_{\text{ref}}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{\text{ref}}) \in \mathbb{R}_+, \quad (8)$$

which is unitarily invariant and fully accounts for complex correlations, and where the following statistical parameters are estimated from $\{\mathbf{z}_n\}_{n=1}^N$:

$$\hat{\boldsymbol{\mu}}_{\text{ref}} = \frac{1}{N} \sum_n \mathbf{z}_n, \quad \hat{\boldsymbol{\Sigma}}_{\text{ref}} = \frac{1}{N-1} \sum_n (\mathbf{z}_n - \hat{\boldsymbol{\mu}}_{\text{ref}})(\mathbf{z}_n - \hat{\boldsymbol{\mu}}_{\text{ref}})^H, \quad (9)$$

and the decision uses a percentile threshold on S_{Maha} calibrated on H_0 to meet PFA α .

For each score, λ is set on an independent H_0 validation set to reach PFA = α : $\lambda = \text{Percentile}(S_{\text{clutter}}, 1 - \alpha)$.

We report Probability of Detection (PD) versus SNR, averaged over multiple Monte-Carlo trials, across Gaussian, compound Gaussian, and real data.

Algorithm 2 Mahalanobis Detection

Require: Encoder Enc , clutter set $\mathcal{D}_0$, test sample $\mathbf{x}$, target PFA α

- 1: Build $\mathbf{z}_i \leftarrow \text{Enc}(\mathbf{x}_i)$ for all $\mathbf{x}_i \in \mathcal{D}_0$
 - 2: Estimate $(\hat{\boldsymbol{\mu}}_{\text{ref}}, \hat{\boldsymbol{\Sigma}}_{\text{ref}})$ via (9)
 - 3: Calibrate $\lambda_{\text{Maha}} = \text{Percentile}(\{S_{\text{Maha}}(\mathbf{z}_i)\}, 1 - \alpha)$
 - 4: On $\mathbf{x}$: get $\mathbf{z}_{\text{test}}$, compute $S_{\text{Maha}}(\mathbf{z}_{\text{test}})$ via (8)
 - 5: Decide H_1 iff $S_{\text{Maha}}(\mathbf{z}_{\text{test}}) > \lambda_{\text{Maha}}$
-

4. RESULTS AND DISCUSSION

We assess Probability of Detection P_d vs. SNR at fixed $P_{fa} = 10^{-2}$ under three conditions: (a) cGN+AWGN, (b) cCGN+AWGN, and (c) real CSIR data. Figures report solid lines for the mean over all cells except Doppler 0, and dashed lines for the Doppler-0 cell.

4.1. Signal and Noise Characteristics

The target echo is modeled as $\alpha = \sqrt{\text{SNR}} \exp(2i\pi\phi) / \sqrt{m}$, where $\phi \in [0, 1]$, and the steering vector is given by $\mathbf{p} = (1, \exp(2i\pi d/m), \dots, \exp(2i\pi d(m-1)/m))^T$ for $m = 16$ bins, where d representing the target normalized Doppler bin index. The clutter covariance matrix follows $\boldsymbol{\Sigma}_c = \mathcal{T}(\rho)$ with $\rho = 0.5$, while texture components τ and τ_k are sampled from a Gamma distribution $\Gamma(\mu, 1/\mu)$ with $\mu = 1$. For adaptive detectors, the covariance estimation uses SCM and Tyler estimators with $K = 2m$ independent secondary samples.

Concerning experimental data, we used the South Africa CSIR (Council for Scientific and Industrial Research) dataset, which comprises radar echoes from sea clutter and naval targets, collected during two campaigns: in 2006 at the Overberg




Fig. 2: Detection performance under different noise configurations, ($P_{fa} = 10^{-2}$).




Fig. 3: On the left: range-slow time map. On the right: Doppler-pulse spectrogram for range bin 5605.

Test Range with the Fynment radar, and in 2007 on Signal Hill with an experimental monopulse radar. It spans diverse waveforms, azimuths, ranges, and environmental conditions (see Fig. 3).

4.2. CVAE Training Configuration

The VAE is trained on clutter-plus-noise Doppler profiles for each noise scenario. The dataset $\mathcal{D}_{H_0}$ consists of 15,000 samples, two-thirds of which are allocated for training and the remaining third for validation. Training is performed over 50 epochs using the Adam optimizer [20] with a learning rate of 10^{-3} . The loss function $\mathcal{L}_{\text{CVAE}}$ incorporates a regularization parameter of $\alpha = 10^2$ for the CVAE-MSE and $\alpha = 10^{-3}$. The latent space dimension is set to 12 for the VAE-MSE and to 32 for the Mahalanobis and KLD latent detectors. Once trained, detection is performed using the reconstruction loss $\mathcal{L}_{\text{rec}}$, with $P_{fa} = 10^{-2}$, determined from a validation set of 5,000 independently generated samples.

4.3. Detection Performance

For cGN+AWGN, CVAE-MSE attains the earliest rise and saturation, yielding the smallest SNR needed to reach a target P_d . Among latent scores, Hermitian Mahalanobis outperforms KLD. ANMF-FP trails the latent Mahalanobis in this Gaussian setting. Doppler-0 behavior is broadly consistent with the mean, with minor shifts.

For cCGN+AWGN, CVAE-MSE remains best overall.

However, ANMF-FP becomes the second-strongest detector under compound Gaussian clutter, surpassing latent Mahalanobis and clearly outperforming KLD. For several detectors (Mahalanobis, KLD), Doppler 0 is slightly *easier* than the mean (left-shifted curves), whereas VAE-MSE exhibits a very slight degradation at Doppler 0.

On experimental CSIR data, ANMF-FP provides the best performance for the Doppler 0, achieving the lowest SNR at moderate-to-high P_d . On the mean over cells, VAE-MSE ranks first, followed by ANMF-FP and Mahalanobis; KLD is the weakest. Crucially, on the Doppler-0 cell, VAE-MSE degrades markedly, while Mahalanobis remains significantly more robust (improved P_d at the same SNR), and ANMF-FP retains the lead. This suggests that (i) classical matched-filtering with Tyler robust covariance estimate (ANMF-FP) handles structured, real clutter particularly well, and (ii) incorporating latent covariance geometry (Mahalanobis) alleviates the specific Doppler-0 failure mode observed with pure reconstruction.

All methods saturate near $P_d \approx 1$ at sufficiently high SNR in synthetic settings, with VAE-MSE reaching saturation first. On real data, ANMF-FP saturates fastest, Mahalanobis catches up steadily, and KLD exhibits delayed growth.

5. CONCLUSION

We compared complex latent CVAE scores (Mahalanobis, KLD) with VAE-MSE and ANMF-FP across synthetic and real radar data at fixed P_{fa} .

On synthetic (cGN/cCGN) data, VAE-MSE is consistently the most SNR-efficient detector; under compound-Gaussian clutter, ANMF-FP is the strongest non-reconstruction alternative, ahead of latent Mahalanobis, while KLD is the weakest. On CSIR data, ANMF-FP is best when considering both the mean and the Doppler-0 cell; if one looks only at the mean (excluding Doppler-0), VAE-MSE slightly leads, but it degrades at Doppler-0 where Mahalanobis is markedly more robust. KLD underperforms in all settings.

6. REFERENCES

- [1] E. J. Kelly, "An adaptive detection algorithm," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-22, no. 2, pp. 115–127, 1986.
- [2] F. C. Robey, D. R. Fuhrmann, E. J. Kelly, and R. Nitzberg, "A CFAR adaptive matched filter detector," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, pp. 208–216, 1992.
- [3] L. L. Scharf and B. Friedlander, "Matched subspace detectors," *IEEE Transactions on Signal Processing*, vol. 42, no. 8, pp. 2146–2157, 1994.
- [4] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *Foundations and Trends® in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.
- [5] Y. A. Rouzoumka, E. Terreaux, C. Morisseau, J.-P. Ovarlez, and C. Ren, "Out-of-distribution radar detection in compound clutter and thermal noise through variational autoencoders," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025.
- [6] Y. Xie, T. Arildsen, and Z.-H. Tan, "Complex recurrent variational autoencoder for speech resynthesis and enhancement," in *International Joint Conference on Neural Networks (IJCNN)*, 2024, pp. 1–7.
- [7] S. N. Marimont and G. Tarroni, "Anomaly detection through latent space restoration using vector quantized variational autoencoders," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 2021, pp. 1764–1767.
- [8] I. Mitiche, A. Salimy, F. Werner, P. Boreham, A. Nesbitt, and G. Morison, "OODCN: Out-of-distribution detection in capsule networks for fault identification," in *2021 29th European Signal Processing Conference (EUSIPCO)*, 2021, pp. 1686–1690.
- [9] Z. Bukhsh and A. Saeed, "On out-of-distribution detection for audio with deep nearest neighbors," in *2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [10] C. Trabelsi, O. Bilaniuk, Y. Zhang, D. Serdyuk, S. Subramanian, J. F. Santos, S. Mehri, N. Rostamzadeh, Y. Bengio, and C. J. Pal, "Deep complex networks," in *International Conference on Learning Representations (ICLR)*, 2018.
- [11] J. A. Barrachina, C. Ren, C. Morisseau, G. Vieillard, and J.-P. Ovarlez, "Comparison Between Equivalent Architectures of Complex-valued and Real-valued Neural Networks - Application on Polarimetric SAR Image Segmentation," *Journal of Signal Processing Systems*, July 2022.
- [12] Y. Hou, Z. Chen, M. Wu, C.-S. Foo, X. Li, and R. M. Shubair, "Mahalanobis distance based adversarial network for anomaly detection," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 3192–3196.
- [13] Y. A. Rouzoumka, E. Terreaux, C. Morisseau, J.-P. Ovarlez, and C. Ren, "Complex-Valued Variational Autoencoders for Radar Detection in Joint Compound Gaussian Clutter and Thermal Noise," *EUSIPCO*, 2025.
- [14] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," Tech. Rep., SNU Data Mining Center, Seoul National University, 2015.
- [15] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [16] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [17] H. H. Nguyen, C. N. Nguyen, X. T. Dao, Q. T. Duong, D. P. Thi Kim, and M.-T. Pham, "Variational autoencoder for anomaly detection: A comparative study," in *ICCE 2024 - IEEE Tenth International Conference on Communications and Electronics*, 2024.
- [18] P. C. Mahalanobis, "On the generalised distance in statistics," *Proceedings of the National Institute of Sciences of India*, vol. 2, pp. 49–55, 1936.
- [19] T. Nakashika, "Complex-valued variational autoencoder: A novel deep generative model for direct representation of complex spectra," in *Interspeech 2020*, 2020, pp. 2002–2006.
- [20] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.