SAR Anomaly Detection With Complex-Valued Variational Autoencoders

X. H. Nguyen*†, J. Fix‡, C. Ren*, J.-P. Ovarlez*† and J. Frontera-Pons†
*SONDRA, CentraleSupélec, Université Paris-Saclay 91190 Gif-sur-Yvette, France,
†DEMR, ONERA, Université Paris-Saclay, 91123 Palaiseau, France,
‡LORIA, CentraleSupélec, Université Paris-Saclay, 57000 Metz, France.

***Abstract*—While recent deep learning methods have shown promising results for SAR imagery and anomaly detection, their performance remains constrained by real-valued architectures that discard phase information. In this paper, we evaluate the performance of a complex-valued Variational AutoEncoder (VAE) for unsupervised SAR anomaly detection. The model learns the SAR anomaly-free data distribution and detects anomalies through reconstruction errors, further enhanced using a covariance matrix–based change detector. Experiments conducted on X-band ONERA polarimetric SAR imagery and synthetic data demonstrate improved anomaly detection performance compared to the classical Reed–Xiaoli detector, particularly in challenging scenarios involving heterogeneous speckle and low signal-to-clutter ratio.**

***Index Terms*—Anomaly detection, SAR Variational AutoEncoder, Complex-Valued Neural Networks**

I. Introduction

Interest in spaceborne and airborne remote sensing has grown substantially, driven by the deployment of numerous satellite missions (e.g., Sentinel series, and more recently Biomass) and the advancement of airborne platforms such as UAVSAR. Among remote sensors, Synthetic Aperture Radar (SAR) is particularly notable because it is an active system that transmits electromagnetic pulses and records the backscattered signals, thereby enabling continuous acquisition regardless of weather conditions and solar illumination. Moreover, the physical properties of commonly employed radar frequencies-such as P-, L-, and X-bands- exhibit distinct penetration behaviors, making SAR uniquely suited to a wide range of applications.

Anomaly detection in SAR imagery has become a critical research area, yet it presents several unique challenges. In multidimensional and complex-valued SAR signals, the presence of speckle fluctuations causes many false alarms. Most importantly, anomalies are unknown in non-controlled environments, rendering annotation difficult and costly. Consequently, despite the abundance of SAR data, labeled examples remain scarce, motivating the adoption of unsupervised methodologies and self-supervised algorithms.

While Deep Learning has shown remarkable progress [1], only a few researchers have tackled the problem of anomaly detection in SAR imaging [2–7]. Most promising outcomes are mainly based on two general methodologies: analyzing reconstruction errors in reconstruction or identifying outliers in a low-dimensional embedding space [8]. A notable commonality among these previous works is their reliance on real-valued neural networks, which consider the magnitude of complex-valued signals only. This inevitably discards important information, since the phase component carries valuable details about SAR diversity—such as polarimetric and interferometric channels, or multi-band and multi-view representations—thus leaving open the question of whether complex-valued architectures could better exploit the inherent properties of SAR data.

Preliminary studies [9–11] have already highlighted the potential benefits of complex-valued neural networks (CVNNs) in handling complex-valued signals; however, comprehensive demonstrations and large-scale applications are still required to establish their practical impact. Joining these efforts, we investigate the application of CVNNs to the anomaly detection problem in SAR imagery by exploiting the Variational AutoEncoder framework [12] in the complex domain [13]. The proposed architecture processes complex-valued inputs, thereby exploiting the phase attribute. This design enhances the quality of latent space representation, particularly in Polarimetric SAR (PolSAR), where phase information encodes essential physical and geometrical properties.

Notations: Matrices are denoted by bold uppercase letters, and vectors by bold lowercase letters. For any matrix $\mathbf{A}$ or vector, $\mathbf{A}^H$ represents the Hermitian transpose of $\mathbf{A}$. The symbols $\odot$ and $\oslash$ denote the Hadamard element-wise product and division, respectively. For vectors, the operator $^\circ$ is applied element-wise (e.g., power, logarithm, modulus).

II. Related works

Statistical anomaly detectors. The mathematical formulation of the anomaly identification problem is presented below. Traditional approaches have relied on the following two binary hypothesis tests:

$$\begin{cases} H_0 : \mathbf{x}_t = \mathbf{c}, & \{\mathbf{x}_i = \mathbf{c}_i\}_{i\in[\![1,n]\!]} \ , \\ H_1 : \mathbf{x}_t = \mathbf{c} + A\,\mathbf{p}, & \{\mathbf{x}_i = \mathbf{c}_i\}_{i\in[\![1,n]\!]} \ . \end{cases} \tag{1}$$

where $\mathbf{x}_t \in \mathbb{C}^4$ denotes a full-polarimetric test pixel, which may correspond either to a clutter pixel $\mathbf{c}$ or to an anomalous pixel $\mathbf{c} + A\,\mathbf{p}$ with unknown magnitude A and steering vector $\mathbf{p}$. All n pixels, $\mathbf{C} = (\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n) \in \mathbb{C}^{4\times n}$, constitute the secondary data used to estimate the background covariance matrix. The classical Generalized Likelihood Ratio framework leads to the design of the well-known Reed-Xiaoli Detector

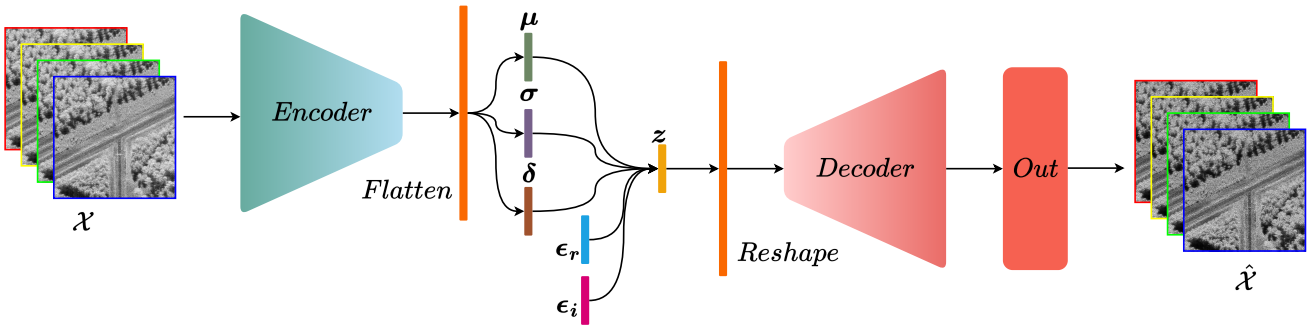

Fig. 1: Architecture of a complex-valued Variational AutoEncoder, where $\mathcal{X}$ and $\hat{\mathcal{X}}$ denote the original Single-Look complex SAR image and its reconstructed version, respectively.

(RXD) [14, 15], which designs the Mahalanobis distance as a point-to-distribution measure:

$$RXD(\mathbf{x}_t) = \mathbf{x}_t^H \, \hat{\mathbf{\Sigma}}_{SCM}^{-1} \, \mathbf{x}_t \underset{H_0}{\overset{H_1}{\gtrless}} \lambda \,, \tag{2}$$

where $\hat{\mathbf{\Sigma}}_{SCM} = \mathbf{C}\mathbf{C}^H/n$ is the Sample Covariance Matrix (SCM) computed on the secondary dataset.

Other classical statistical detectors defined in [16] and [17] also exploit background covariance estimation to discriminate anomalies or targets from clutter. However, the assumption of Gaussianity increases their sensitivity to clutter heterogeneity and limits their ability to capture complex spatial and contextual information. Other alternatives have alleviated these constraints, notably the Tyler estimator [18, 19, 15], which provides a robust distribution-free estimation of the covariance matrix, allowing effective modeling of heavy-tailed and non-Gaussian backgrounds commonly encountered in SAR imagery.

Deep learning reconstruction-based anomaly detector. Recent efforts [4, 5, 2, 3] have established new benchmarks in SAR anomaly detection by leveraging the rapid progress of Deep Learning. In particular, Muzeau et.al. [2, 3] and Sinha et.al. [5] employed variants of deep convolutional AutoEncoders to model background clutter distributions and quantify anomalies through reconstruction errors, with Muzeau further demonstrating that the inclusion of a despeckling step can significantly reduce the probability of false alarms. In parallel, Mabu [4] proposed an unsupervised Adversarial AutoEncoder combined with a one-class SVM to identify anomalous instances in the latent space. More recent works by Chauvin [6] and Ibarra [7] extended this line of research by introducing anomaly detection on public SAR datasets within the patch distribution modeling framework PaDiM [20].

III. Anomaly Detection in Complex-Valued SAR Imaging

We leverage a variational autoencoder (VAE) extended to the complex domain to process complex-valued data. Traditionally, two common approaches are used to train real-valued neural networks on complex-valued inputs. The first considers only the magnitude of the complex signals, thereby discarding phase information and resulting in a lossy representation. The second, more pragmatic, approach stacks the real and imaginary parts as separate channels; however, this may break their underlying algebraic relationships and significantly distort the original phase information during deep operations such as convolutions or normalization. In contrast, employing dedicated complex-valued architectures enables efficient processing while preserving the intrinsic magnitude–phase relationships of the data [9–11].

A. Complex-Valued Variational AutoEncoder

Initially introduced by [12], VAEs belong to the family of generative neural networks based on Deep Learning. They aim to approximate the data distribution using a variational Bayesian approach and can be seen equivalently as a regularized AutoEncoder. Tailoring VAE to the complex domain requires learning from potential non-circular data characterized by the covariance and pseudo-covariance matrices in the Gaussian case. Several works [21, 13] have already studied this adaptation; we thereby adopt their strategy to learn the intrinsic representations of Single-Look Complex SAR data. The overall complex-valued VAE architecture is illustrated in Figure 1.

In real-valued VAEs, the latent variables $\boldsymbol{z}$ are typically assumed to follow a standard Normal *a priori* distribution. By analogy, in complex-valued VAEs, the latent vector $\boldsymbol{z}$ is modeled using a complex circular Normal distribution $\mathcal{CN}(\mathbf{0}, \mathbf{I})$.

Complex-valued Encoder & Decoder: Let $\mathcal{X} \in \mathbb{C}^{4\times h\times w}$ be the four polarization channels of complex-valued SAR with spatial dimensions height h and width w. Let $\hat{\mathcal{X}} \in \mathbb{C}^{4\times h\times w}$ denote its reconstruction delivered by the complex-valued VAE. In autoencoder-based architectures, the encoder spatially compresses $\mathcal{X}$ into a latent representation $\mathbf{z} \in \mathbb{C}^q$. Each encoding block consists of two stacked complex-valued $Convolution$–$BatchNorm$–$Activation$ sequences, where the activation function is the phase-preserving $\mathrm{modReLU}$ introduced in [22]:

$$\mathrm{modReLU}(x) = \mathrm{ReLU}(|x| + b)\, e^{j\,\arg(x)} \,, \tag{3}$$

For downsampling, instead of using Average Pooling or magnitude-based Max Pooling operation, we adopt strided convolutions for the first convolutional layer of each encoder block, allowing the network to learn adaptive feature extraction

while better preserving the complex spatial structures of SAR images.

Unlike real-valued VAEs, the encoder outputs three parameter vectors for the latent space of dimension q: the mean $\boldsymbol{\mu} \in \mathbb{C}^q$, the covariance $\boldsymbol{\sigma}^{\circ 2} \in \mathbb{R}^{q+}$, and the pseudo-covariance $\boldsymbol{\delta} \in \mathbb{C}^q$. The latter provides additional flexibility for accurately modeling the second-order structure of complex-valued input data.

The decoder reconstructs the input image by mirroring the encoder architecture and employing two stacked $Convolution$–$BatchNorm$–$modReLU$ layers per block. Upsampling is performed using strided complex transposed convolution layers. The final reconstructed image $\hat{\mathcal{X}}$ is obtained by applying a convolutional layer to the output of the last decoder block.

Complex-valued reparameterization trick. We adopt the reparameterization trick for complex-valued VAE, originally introduced by Nakashika [13] and further extended in [21]. Specifically, the three complex-valued parameters $\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\delta}$ are sampled as follows:

$$\boldsymbol{z} = \boldsymbol{\mu} + \boldsymbol{k_r} \odot \boldsymbol{\epsilon}_r + i\, \boldsymbol{k_i} \odot \boldsymbol{\epsilon}_i \,, \tag{4}$$

where $\begin{cases} \boldsymbol{k}_r = 2^{-\frac{1}{2}} (\boldsymbol{\mu} + \boldsymbol{\delta}) \oslash (\boldsymbol{\sigma} + \mathrm{Re}(\boldsymbol{\delta}))^{\circ \frac{1}{2}} \,, \\ \boldsymbol{k}_i = 2^{-\frac{1}{2}} \left(\boldsymbol{\sigma}^{\circ 2} - (|\boldsymbol{\delta}|^{\circ})^{\circ 2}\right)^{\circ \frac{1}{2}} \oslash (\boldsymbol{\sigma} + \mathrm{Re}(\boldsymbol{\delta}))^{\circ \frac{1}{2}} \end{cases}$,

and where $\boldsymbol{\epsilon}_r \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$. The pseudo-covariance is computed with $\boldsymbol{\delta} = \boldsymbol{\alpha} \odot \boldsymbol{\sigma}$ with $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_q)^T \in \mathbb{C}^q, |\alpha_i| < 1, \forall i \in [\![1, q]\!]$.

Loss function. Traditionally, the VAE formalism [12] leads to minimize the *Evidence Lower BOund* loss function:

$$\mathrm{ELBO} = \mathcal{L}_{rec} - D_{KL}(q(\mathbf{z}|\mathcal{X}) \| p(\mathbf{z})) \,. \tag{5}$$

where $L_{rec} = \mathbb{E}_{q(\mathbf{z}|\mathcal{X})} [\log (p(\mathcal{X}|\mathbf{z})]$ with $p(\mathcal{X}|\mathbf{z})$ being the generative distribution. This variational bound to be minimized consists of a reconstruction term, which ensures high-quality image reconstruction, and a regularization term, which enforces a closeness between the latent space distribution and the desired prior. For a complex-valued Gaussian model, the reconstruction loss $\mathcal{L}_{rec}$ is formulated as an L^2 loss, while the Kullback-Leibler (KL) divergence [21] is expressed as:

$$\begin{aligned} D_{KL}(q(\mathbf{z}|\mathcal{X}) || p(\mathbf{z})) \quad &= \quad ||\boldsymbol{\mu}||^2 + \\ &\mathbf{1}_q^T \left(\boldsymbol{\sigma} - 1 - \frac{1}{2} \log^{\circ} \left(\boldsymbol{\sigma}^{\circ 2} - (|\boldsymbol{\delta}|^{\circ})^{\circ 2} \right) \right) . \end{aligned} \tag{6}$$

Here we consider the β-VAE [23], in which the importance of the regularization term is weighted by the coefficient β. The new loss function $\mathcal{L}_{VAE}$ to be minimized becomes:

$$\mathcal{L}_{VAE} = \mathcal{L}_{rec} + \beta\, D_{KL} \,. \tag{7}$$

B. Change detection

SAR images, which exhibit high spatial and spectral dynamics, therefore require statistics to be calculated locally. We complete the anomaly detection process with a complex-valued VAE by analyzing the reconstruction error between the input image and its reconstructed version with covariance. For each pixel $k \in [\![0, h-1]\!], l \in [\![0, w-1]\!]$, we define a boxcar $\mathcal{B}_{k,l}$ of cardinal $|\mathcal{B}_{k,l}|$, surrounding and centered on the pixel k, l. An anomaly score is defined as the Frobenius norm between the covariance matrices estimated on a boxcar from the reconstructed signal and from the original signal:

$$A_{k,l} = \left\| \hat{\boldsymbol{\Sigma}}^{\hat{\mathcal{X}}}_{k,l} - \hat{\boldsymbol{\Sigma}}^{\mathcal{X}}_{k,l} \right\|_F^2 , \tag{8}$$

where $\hat{\boldsymbol{\Sigma}}_{k,l} \in \mathbb{C}^{4\times 4}$ is the covariance matrix of the polarimetric channels. For each $\mathcal{B}_{k,l}$ of $\mathcal{X}$ (and similarly for each $\mathcal{B}_{k,l}$ of $\hat{\mathcal{X}}$), it is calculated as follow:

$$\hat{\boldsymbol{\Sigma}}^{\mathcal{X}}_{k,l} = \frac{1}{|\mathcal{B}_{k,l}|} \sum_{i \in \mathcal{B}_{k,l}} \mathbf{x}_i \, \mathbf{x}_i^H \,. \tag{9}$$

Other change detection methods, based on covariance-equality tests [24], will be investigated in the future.

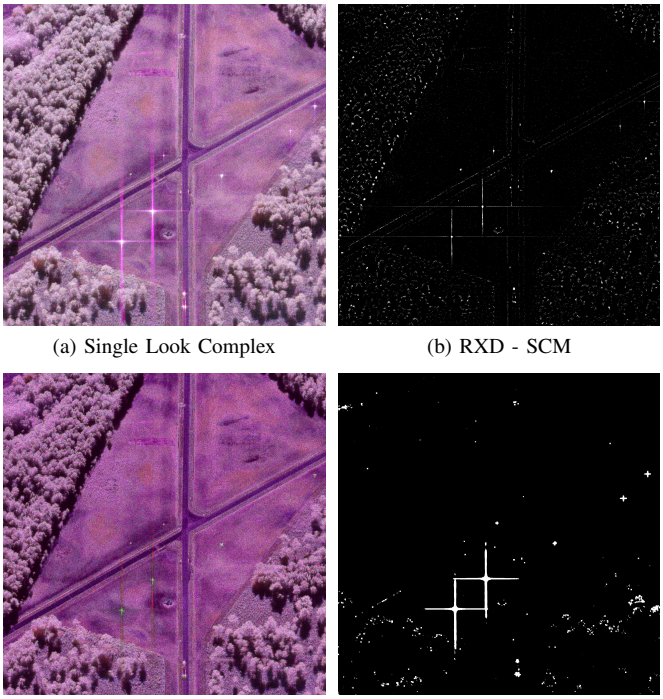

(a) Single Look Complex (b) RXD - SCM

(c) VAE's prediction (d) Frobenius 9×9

Fig. 2: Qualitative comparison between the RXD detector built with SCM and the complex-valued VAE with a Frobenius-based change detector. Images (b) and (d) are displayed using the top 2% of values.

IV. Experimental results

A. Data

We train the proposed complex-valued VAE on a single-look complex (SLC) fully polarimetric SAR image acquired by ONERA's SETHI airborne sensor [25]. The original image has spatial dimensions of 4800×30000, with a pixel resolution of 20 cm in both azimuth and slant-range directions. During training, the image is divided into multiple slightly overlapping patches of size 64×64, using a stride of 50 pixels. We

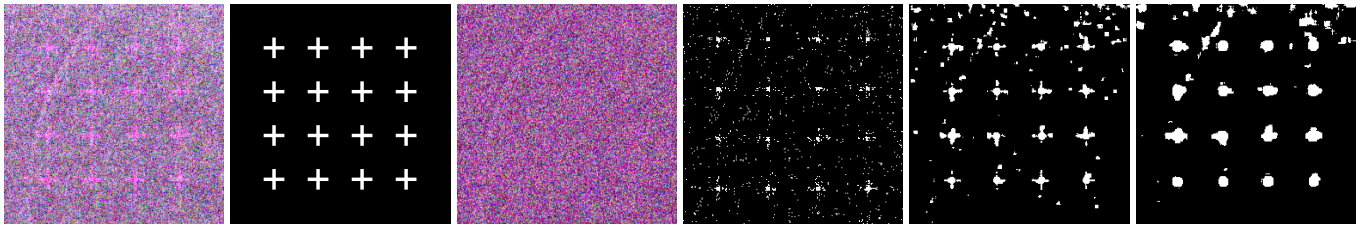

(a) Original (b) Ground truth (c) VAE's prediction (d) RXD SCM (e) Frobenius 5×5 (f) Frobenius 9×9

Fig. 3: Anomaly maps generated by the RXD-SCM detector and the complex-valued VAE with a Frobenius-based change detector. From left to right: input image with synthetic anomalies at an SNR of 5 dB added to homogeneous clutter, ground-truth labels, VAE reconstruction, anomaly map computed using RXD-SCM, and anomaly maps obtained using the Frobenius distance with boxcar sizes of 5×5 and 9×9 for PFA=0.02.

use 80% of the image for training and 20% for validation. For visualization purposes, considering that HV = VH (by principle of reciprocity), all images in this paper are displayed using only the (HH, HV, VV) polarimetric channels mapped to the red, green, and blue colors, respectively.

In SAR imagery, anomalies are generally unknown *a priori*, complicating annotation and quantitative evaluation. To address this limitation, we evaluate the proposed method using synthetic anomalies (see Fig. 3) injected into homogeneous clutter regions at a fixed signal-to-clutter ratio (SNR). The approximate SNR for these synthetic anomalies is defined as:

$$SNR(\mathbf{x}_t) = A^2\, \mathbf{p}^H\, \hat{\mathbf{\Sigma}}_{SCM}^{-1}\, \mathbf{p}\,, \tag{10}$$

where $\mathbf{p} = (1, 0, 0, 1)^T$ is the steering vector corresponding to the co-polarized HH and VV channels, and $\hat{\mathbf{\Sigma}}_{SCM}$ denotes the estimated covariance matrix of the homogeneous clutter.

B. VAE training settings

Our VAE is based on the complex-valued autoencoder proposed in [11], with a modified latent space and the complex-valued reparameterization described in Section III-A. Let `Ck` and `CDk` denote encoder and decoder blocks with `k` filters, respectively. The encoder and decoder architectures consist of `C32 – C64 – C128 – C256` and `CD256 – CD128 – CD64 – CD32`, respectively. All convolutional kernels have a size of 3×3. To better capture the global structure and variability of the input training patches, the convolutional kernel of `C32` is increased to 9×9. All `C32`, `CD32`, `CD64`, `CD128` and `CD256` use a stride of 1.

For an input of size $64 \times 64 \times 4$, the encoder outputs a feature map of size $8 \times 8 \times 256$. In our experiments, compressing the latent space by a factor of 2 yields the best AUC. To improve image reconstruction quality, β is set to 0 during the first five training epochs, then increased linearly from the 6th to 15th epoch up to 10^{-6} and kept constant thereafter. The value of β is heuristically chosen to reach a balance between the reconstruction loss and the KL divergence. We use the Adam optimizer with a learning rate of 10^{-4} and a batch size is set to 320. Training is performed on an NVIDIA RTX 4090 GPU, with an average duration of approximately 3 min 15 seconds. The best model is defined upon the best validation reconstruction loss, which converges after 109 epochs.

C. Anomaly detection performance

For the RXD detector, the Sample Covariance Matrix is estimated using a boxcar of size 31×31 containing an exclusion window (cell guard) of size 21×21.

Fig. 3a shows cross-shaped anomalies with an SNR of 5 dB, which are particularly challenging to detect. As illustrated in Fig. 3c, the VAE output effectively suppresses these anomalies, which is beneficial for subsequent change detection. The detection maps shown in Fig. 3d–f are thresholded at a probability of false alarm (PFA) of 2%. The qualitative results indicate that the complex-valued VAE outperforms the Reed–Xiaoli detector. It is also worth noting that increasing the size of the change-detection boxcar improves detection performance.

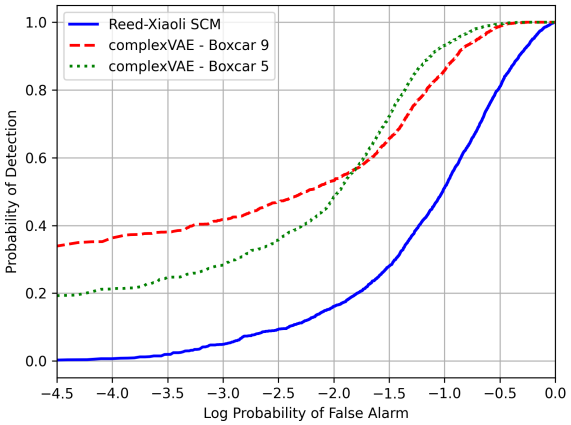

Fig. 4: Probability of Detection versus $\log_{10}(P_{fa})$ for an anomaly with an SNR=5 dB.

V. Conclusion

This paper improves anomaly detection in SAR imagery by extending conventional reconstruction-error analysis with a complex-valued variational autoencoder combined with a covariance change detector. Both qualitative and quantitative evaluations demonstrate superior performance compared to the classical Reed–Xiaoli detector at an equivalent probability of false alarm. Future work will focus on improved learning strategies and advanced statistical detection schemes based on covariance-equality tests to further enhance robustness and detection accuracy.

REFERENCES

[1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[2] M. Muzeau, C. Ren, S. Angelliaume, M. Datcu, and J.-P. Ovarlez, "Self-supervised learning based anomaly detection in synthetic aperture radar imaging," *IEEE Open Journal of Signal Processing*, vol. 3, pp. 440–449, 2022.

[3] M. Muzeau, C. Ren, S. Angelliaume, M. Datcu, and J.-P. Ovarlez, "Self-supervised SAR anomaly detection guided with RX detector," in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, 2023, pp. 1918–1921.

[4] S. Mabu, S. Hirata, and T. Kuremoto, "Anomaly detection using convolutional adversarial autoencoder and one-class SVM for landslide area detection from synthetic aperture radar images," *Journal of Robotics, Networking and Artificial Life*, vol. 8, no. 2, pp. 139–144, 2021.

[5] S. Sinha, S. Giffard-Roisin, F. Karbou, M. Deschatres, A. Karas, N. Eckert, C. Coléou, and C. Monteleoni, "Variational autoencoder anomaly-detection of avalanche deposits in satellite sar imagery," in *Proceedings of the 10th International Conference on Climate Informatics*, 2020, pp. 113–119.

[6] L. Chauvin, S. Gupta, A. Ibarra, and J. Peeples, "Benchmarking suite for synthetic aperture radar imagery anomaly detection (SARIAD) algorithms," in *Algorithms for Synthetic Aperture Radar Imagery XXXII*. SPIE, 2025, vol. 13456, pp. 95–101.

[7] A. Ibarra and J. Peeples, "Patch distribution modeling framework adaptive cosine estimator (PaDiM-ACE) for anomaly detection and localization in synthetic aperture radar imagery," in *Algorithms for Synthetic Aperture Radar Imagery XXXII*. SPIE, 2025, vol. 13456, pp. 102–108.

[8] I. Golan and R. El-Yaniv, "Deep anomaly detection using geometric transformations," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2018, NIPS'18, p. 9781–9791, Curran Associates Inc.

[9] 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*, vol. 95, no. 1, pp. 57–66, 2023.

[10] A. Hirose, *Complex-valued neural networks*, Springer, 2006.

[11] Q. Gabot, J. Fix, J. Frontera-Pons, C. Ren, and J.-P. Ovarlez, "Preserving polarimetric properties in polsar image reconstruction through complex-valued auto-encoders," in *2024 International Radar Conference (RADAR)*. IEEE, 2024, pp. 1–6.

[12] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.

[13] T. Nakashika, "Complex-valued variational autoencoder: A novel deep generative model for direct representation of complex spectra.," in *Interspeech*, 2020, pp. 2002–2006.

[14] Xiaoli Y. and I. S. Reed, "Adaptive detection of signals with linear feature mappings and representations," *IEEE Transactions on Signal Processing*, vol. 43, no. 12, pp. 2953–2963, 1995.

[15] J. Frontera-Pons, M. Veganzones, F. Pascal, and J.-. Ovarlez, "Hyperspectral anomaly detectors using robust estimators," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 9, no. 2, pp. 720–731, 2015.

[16] E. J. Kelly, "An adaptive detection algorithm," *IEEE Transactions on Aerospace and Electronic Systems*, , no. 2, pp. 115–127, 2007.

[17] D. R. Fuhrmann, E. J. Kelly, and R. Nitzberg, "A CFAR adaptive matched filter detector," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, no. 1, pp. 208–216, 1992.

[18] D. E. Tyler, "A distribution-free M-estimator of multivariate scatter," *The Annals of Statistics*, pp. 234–251, 1987.

[19] J. Frontera-Pons, M. Veganzones, S. Velasco-Forero, F. Pascal, J.-P. Ovarlez, and J. Chanussot, "Robust anomaly detection in hyperspectral imaging," in *2014 IEEE Geoscience and Remote Sensing Symposium*. IEEE, 2014, pp. 4604–4607.

[20] T. Defard, A. Setkov, A. Loesch, and R. Audigier, "PaDiM: a patch distribution modeling framework for anomaly detection and localization," in *International Conference on Pattern Recognition*. Springer, 2021, pp. 475–489.

[21] 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," in *33rd European Signal Processing Conference (EUSIPCO)*, 2025.

[22] M. Arjovsky, A. Shah, and Y. Bengio, "Unitary evolution recurrent neural networks," in *International conference on machine learning*. PMLR, 2016, pp. 1120–1128.

[23] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-VAE: Learning basic visual concepts with a constrained variational framework," in *International conference on learning representations*, 2017.

[24] A. Mian, A. Collas, A. Breloy, G. Ginolhac, and J.-P. Ovarlez, "Robust low-rank change detection for multivariate SAR image time series," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 3545–3556, 2020.

[25] R. Baqué, P. Dreuillet, and H. Oriot, "Sethi: Review of 10 years of development and experimentation of the remote sensing platform," in *2019 International Radar Conference (RADAR)*. IEEE, 2019, pp. 1–5.