

RFI Mitigation on Raw SAR Data Using U-Net for Enhanced Signal Reconstruction

Nermine Hendy¹, Akram Al-Hourani¹, Mikolaj Czerkawski², Nicolas Longép  ², Thomas Kraus³, Maximilian Schandri³, Markus Bachmann³, and Haytham M Fayek⁴

¹School of Engineering, STEM College, RMIT University

²-lab Explore Office, ESRIN, European Space Agency (ESA)

³German Aerospace Center (DLR)

⁴School of Computing Technologies, STEM College, RMIT University

January 31, 2025

RFI Mitigation on Raw SAR Data Using U-Net for Enhanced Signal Reconstruction

Nermine Hendy^{1*}, Akram Al-Hourani^{1*}, Mikolaj Czerkawski², Nicolas Longép  ²,
Thomas Kraus³, Maximilian Schandri³, Markus Bachmann³, and Haytham M. Fayek⁴

¹ School of Engineering, STEM College, RMIT University, Melbourne, VIC, Australia.

² Φ-lab Explore Office, ESRIN, European Space Agency (ESA), Frascati, Italy.

³ German Aerospace Center (DLR), M  nchener Str. 20, 82234, Oberpfaffenhofen, Germany.

⁴ School of Computing Technologies, STEM College, RMIT University, Melbourne, VIC, Australia.

* Correspondence: nerminahendy@ieee.org || akram.hourani@rmit.edu.au

Abstract—Synthetic Aperture Radar (SAR) systems are increasingly at risk of Radio Frequency Interference (RFI), driving the need for effective mitigation techniques to maintain data quality in challenging operational settings. This paper introduces a lightweight and effective framework designed to mitigate the impact of RFI directly on raw Radio Frequency (RF) SAR data. Given the scarcity of publicly labeled raw SAR data with interference events, a synthetic RFI dataset is created using a realistic RF-level SAR emulation tailored for spaceborne scenarios. A twin U-Net model is developed to mitigate RFI in affected segments, inferring clean SAR data from contaminated RFI inputs. The model’s mitigation performance is quantitatively validated on unseen synthetic data. This approach has the potential to enhance the quality of SAR imagery impacted by RFI with the advantage of operating directly in the raw signal domain.

Index Terms—Radio frequency interference, RFI mitigation, synthetic aperture radar.

I. INTRODUCTION

Spaceborne Synthetic Aperture Radar (SAR) is an essential tool for high-resolution Earth observation independent of light and weather conditions. SAR provides valuable data for a diverse range of applications, including environmental monitoring, disaster response, and defense operations. However, due to the inherent use of Radio Frequency (RF) spectrum, SAR is increasingly susceptible to Radio Frequency Interference (RFI), posing significant challenges to the reliability of SAR data. RFI, caused by other transmitters overlapping with the SAR operating band, corrupts the data and introduces impairing artifacts. A common form of RFI is Narrow-Band Interference (NBI), characterized by its narrow Power Spectral Density (PSD) and often appears as saturated lines that obscure critical features in the data. Mitigating such unexpected RFI is challenging due to its unpredictable location and nature, rendering traditional filtering techniques insufficient for effective mitigation. Addressing these challenges is essential for maintaining the integrity and utility of SAR imagery.

Deep Learning (DL) offers transformative solutions to complex RFI challenges in SAR applications, leveraging architectures such as deep residual networks [1], Generative Adversarial Network (GAN)-based frameworks [2], DL-based filtering algorithms [3], and deep segmentation networks [4], [5] (also

explored in other radar signal processing domains [6]). While these models outperform traditional methods, their application to Time-Frequency Domain (TFD) data involves significant preprocessing and computational demands, limiting real-world scalability. Hybrid approaches like Prior-induced Interference Suppression Network (PISNet) [7] combine DL with semi-parametric techniques, achieving strong results at the cost of extensive tuning. Feature decomposition models such as Multi-dimensional Calibration and Suppression Network (MCSNet) [8] and Combined-Attention Restoration Network (CARNet) [9] enhance RFI suppression but depend on reference images and focus on processed rather than raw SAR data, limiting practical utility. Detection-focused models like Canny Edge Detector and Vision Transformer (CEVIT) [10] address raw data but overlook mitigation, while [11] underscores the challenges posed by the reliance on large labeled datasets.

Although some work [4] has addressed the need for datasets by introducing Multi-scale Interference Detection (MID) (for RFI detection) and Polygonal-Annotated Interference Segmentation (PAIS) (for RFI segmentation), which consist of Single-Look Complex (SLC) products and corresponding Ground Range Detected (GRD) quick-looks from Sentinel-1 [12], a significant gap persists: the absence of publicly available labeled datasets specifically for raw RFI SAR data. This limitation highlights the pressing need for such datasets along with methods capable of directly process real contaminated raw SAR data without relying on extensive preprocessing or manually labeled datasets.

This paper introduces a twin U-shaped Network (U-Net)-based model designed to directly mitigate RFI from raw SAR data, eliminating the need for a detection step or hand-crafted features. The proposed model is trained on emulated raw SAR data synthesized using a realistic spaceborne emulator. By generating data at the RF level, full control is maintained over the injected RFI signals, ensuring precise ground truth data for training. The model’s performance is rigorously evaluated on an unseen synthetic test dataset, demonstrating its ability to achieve superior interference mitigation while preserving data fidelity. The primary contributions of this work include: (i) a systematic approach for generating synthetic datasets tailored




Fig. 1: The RFI mitigation framework utilizing the proposed twin-U-Net architecture. The diagram illustrates the pre-processing steps for both magnitude and phase components, including scaling, wrapping, and down-sampling, before passing the processed data to the respective U-Net models. The outputs are combined to reconstruct the raw predicted clean data.

for DL-based RFI mitigation, (ii) the development of a twin U-Net architecture for direct RFI mitigation from raw SAR data, and (iii) a quantitative validation using the simulated data.

II. METHODOLOGY

Similar to conventional Machine Learning (ML) applications, the proposed framework operates in two phases (i) a *training phase* using labeled synthetic data (generated automatically with the simulation engine), and (ii) a *testing phase* involving unseen labeled synthetic data, as illustrated in Fig. 1. Dataset preparation and the model architecture are detailed below.

A. Raw Synthetic Dataset

The synthetic dataset was generated from emulated raw SAR data using **SAR EMULATOR** for **Spaceborne Applications** (SEMUS) [13], an RF-level emulator developed to replicate realistic spaceborne scenarios. To enhance realism, the dataset generation targeted regions with historically high RFI occurrences, identified through European Space Agency (ESA) RFI maps along with urban and coastal hotspots prone to interference [14].

For this experiment, parameters from a real, commonly used sensor of Sentinel-1 were employed to create realistic synthetic data for selected 612 cities globally, with each city imaged in a single swath scan. To optimize computational efficiency, scans were limited to one second, producing 2D raw SAR data with dimensions of 1871×5565 pixels in azimuth and range directions. Diverse RFI scenarios were simulated by introducing NBI, modeled as randomized Linear Frequency Modulated (LFM) chirps, based on realistic interference patterns observed in TerraSAR-X [15] and Sentinel-1 [16]. For further details on generating the RFI, please refer to the research in [5]. Each clean sample was paired with eight RFI-contaminated variations, yielding 4,896 total

TABLE I: EMULATION PARAMETERS [5]

Parameter	Symbol	Value
Eccentricity	e	7.9928×10^{-4}
Argument of periaapsis	ω	90°
Inclination	i	98.176°
Right ascension of the ascending node	Ω	$[-180^\circ, 0^\circ]$
True anomaly	ν	$[-180^\circ, 0^\circ]$
Radar centre frequency	f_o	5.405 GHz
Radar bandwidth	B	100 MHz
RFI center frequency shift	Δf	$[-0.01f_o, 0.01f_o]$
RFI bandwidth	B_I	$[0.1B, 0.3B]$
RFI duty cycle	D_I	$[70\%, 90\%]$
RFI Pulse repetition frequency	PRF_I	$[1, 5]$ kHz
RFI pulse width	T_{pI}	D_I/PRF_I
RFI chirp rate	β_I	B_I/T_{pI}
RFI transmitter latitude shift	$\Delta\theta_{\text{LAT}}$	$[-0.03^\circ, 0.03^\circ]$
RFI transmitter longitude shift	$\Delta\theta_{\text{LON}}$	$[-0.02^\circ, 0.02^\circ]$
RFI transmitter gain (Isotropic)	G_I	0 dB
PSD-SIR	γ	$[-5, 5]$ dB

samples. The orbital configurations, SAR parameters, and RFI characteristics are detailed in Table I.

The dataset comprises both magnitude and phase components saved in the quadrature format $s_{qd} = A \cdot e^{j\phi}$, where A and ϕ represent the contaminated magnitude and phase components, respectively. The magnitude data was scaled by a factor of 1×10^{13} to align its amplitude range with real SAR data of Sentinel-1. In contrast, the phase data, bounded within $[-2\pi, 2\pi]$, was wrapped to the range $[-\pi, \pi]$ using modulo arithmetic and normalized by dividing by π to ensure consistency and stability during training.

Furthermore, to reduce the computational burden in this prototype study, both magnitude and phase arrays are downsampled to a shape of 256×256 pixels using bilinear interpolation (which is performed after the magnitude is scaled and the phase is wrapped). The dataset was randomly split into 400 cities (3,600 samples) for training, 50 cities (400 samples) for validation, and 162 cities (1,296 samples) for testing. This split ensures that cities used for training, along with all their RFI variants, are excluded from both validation and testing datasets. Contaminated samples were used as model input, while clean samples served as ground truth in the supervised setting.

B. Model Architecture

The proposed approach employs twin U-Nets, independently processing the magnitude and phase components. Since RFI manifests differently on each component, this separation prevents phase channel features from dominating magnitude predictions, ensuring balanced and effective model performance. Consequently, the magnitude U-Net predicts clean amplitude values, while the phase U-Net refines the phase information. The outputs are combined to reconstruct the clean signal using $\hat{s}_{qd} = \hat{A} \cdot e^{j\hat{\phi}}$, where $\hat{A}$ and $\hat{\phi}$ represent the predicted clean magnitude and phase components, respectively.

The utilized U-Net architecture incorporates multiple down-sampling stages to extract features from the signal at different scales. Each stage employs ConvNeXt blocks [17] for feature

TABLE II: ARCHITECTURE OF THE PROPOSED U-NET MODEL FOR SAR RFI MITIGATION

Layer Name	Details	Purpose
Downsampling Path		
ConvNeXt Block 1	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 1	Initial mapping of input data
ConvNeXt Block 2	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 64; Residual connection	Extracts and refines features
Residual Block 1	LinearAttention + LayerNorm	Attention mechanism
Downsample Block 1	Conv2d: kernel size = 4x4, stride = 2, padding = 1	Reduces spatial dimensions
ConvNeXt Block 3	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 128; Residual connection	Captures deeper features
Residual Block 2	LinearAttention + LayerNorm	Refines feature representation
Downsample Block 2	Conv2d: kernel size = 4x4, stride = 2, padding = 1	Further reduces spatial dimension
ConvNeXt Block 4	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 256; Residual connection	Extracts complex hierarchical features
Residual Block 3	LinearAttention + LayerNorm	Attention mechanism
Middle		
Bottleneck Block	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 512	Bottleneck for processing features
Upsampling Path		
Upsample Block 1	ConvTranspose2d: kernel size = 4x4, stride = 2, padding = 1	Upscales spatial dimensions
ConvNeXt Block 5	Conv2d: kernel size = 7x7, stride = 1, padding = 3, groups = 64; Residual connection	Final refinement of features
Final Conv Layer	Conv2d: kernel size = 1x1	Reduces channels to match output size

extraction process that are subsequently injected into the corresponding upsampling stages, allowing uninterrupted information flow from the input signal to the reconstructed output. The implementation of skip connections bridges the encoder and decoder paths, preserving critical spatial information and enabling efficient feature reuse. Linear attention is incorporated to capture long-range dependencies in the input data, allowing the network to effectively identify and mitigate interference patterns that may be spatially or temporally distant within the contaminated SAR signal. The output layer generates a single-channel prediction map of the clean magnitude and phase data. The U-Net architecture is summarized in Table II, explaining the downsampling path, middle bottleneck, and upsampling path, clearly showcasing the role of each block within the network.

The network is optimized to minimize pixel-wise differences between the predicted and clean data using Mean Square Error (MSE) loss, computed separately for magnitude and phase, with the combined metric given as $MSE = \frac{1}{2} (MSE_{\text{magnitude}} + MSE_{\text{phase}})$. Coefficient of Determination (R^2) score assesses how well the predictions capture variance in the clean data for both components, ensuring reliable performance across magnitude and phase.

III. RESULTS AND VALIDATION

A. Experimental Setup

Training was conducted on a system equipped with an Intel Core i9-12900 CPU and an NVIDIA GeForce RTX 3060 GPU. The model was optimized using the Adam optimizer with a learning rate of 1×10^{-4} for a maximum of 50 epochs with a batch size of two. Early stopping with a patience of 20 epochs based on the validation MSE was employed, along with gradient clipping (value of 1.0) and mixed-precision training to enhance stability and ensure robust convergence.

TABLE III: TEST SET EVALUATION METRICS

Metric	Contaminated	After RFI Mitigation
↓Raw Mag MSE	1.413	0.033
↓Masked	1.545	0.028
↓Raw Phase MSE	0.000314	0.000096
↓Masked	0.000355	0.000098
↓Focused MSE	0.00238	0.00051
↑Focused SSIM	0.78324	0.90889

B. Quantitative Analysis on Synthetic Data

Evaluation has been performed by applying the model saved with a checkpoint corresponding to the lowest MSE validation loss. Table III demonstrates the added value of applying the trained mitigation model on unseen data. The first two rows report the MSE associated with the raw magnitude representation, including errors within the masked regions indicating interference. The results reveal a substantial improvement in quality, with the MSE dropping from over 1.413 in the contaminated data to 0.033 after applying the mitigation model. This effect is even stronger over the Masked Interference region.

The next two rows detail the MSE reduction for the raw phase signal. Once again, the mitigation model effectively reduces the error, further validating its capability to reconstruct the clean signal. Finally, the last two rows summarize the MSE and Structural Similarity Index (SSIM) metrics over the focused SAR image formed using the magnitude and phase components. The MSE is notably reduced, and the SSIM metric increases from 0.78324 for the contaminated data to 0.90889 after mitigation, indicating enhanced structural similarity to the clean data.

This improvement in visual quality, as reflected by the increased SSIM, aligns with the qualitative observations from the formed images in Figure 2. The first row (a) shows the contaminated raw data magnitude, while the second row (b) presents the corresponding formed images with interference.




Fig. 2: Visualization of the RFI mitigation process. (a) Raw contaminated data magnitude, (b) corresponding focused contaminated data, (c) raw filtered (predicted) magnitude data, and (d) focused filtered (predicted) data. The images were formed using the RDA algorithm. Each column corresponds to a distinct test sample, representing different urban environments.

The last two rows (c-d) depict the same data after applying the mitigation model. While interference is not entirely eliminated, it becomes subtle and difficult to detect visually, even in cases with high levels of contamination in the input data.

IV. CONCLUSION

This study proposed a robust framework for RFI mitigation in SAR systems, by training a twin U-Net model to process both the magnitude and phase of raw SAR data. To address the lack of real-world labeled RFI-SAR datasets, the framework employed controlled training using synthetic emulated data, which allows to generate virtually unlimited amounts of training data for the task of interference removal. Trained on a single GPU for only 50 epochs, the model gains the capability to generalise to previously unseen data in the test dataset and was consequently proven to act as an effective tool for interference removal, as demonstrated by the improved error

and structural similarity metrics reported in the quantitative analysis.

This work lays the foundation for building interference mitigation models that can automate the process of cleaning the contaminated data directly in the raw format. While the context here focused on the simulated data, it is expected that the next steps of this work will explore how the same model could be applied in a real-world setting to contaminated SAR data acquired by physical sensors.

REFERENCES

- [1] W. Fan, F. Zhou, M. Tao, X. Bai, P. Rong, S. Yang, and T. Tian, "Interference mitigation for synthetic aperture radar based on deep residual network," *Remote Sensing*, vol. 11, no. 14, p. 1654, 2019.
- [2] X. Xu, W. Fan, F. Zhou, and Y. Sun, "RFISM-Net: an interference segmentation and mitigation network for SAR," in *IET International Radar Conference (IRC 2023)*, vol. 2023. IET, 2023, pp. 3402–3408.
- [3] F. Fang, D. Dai, and A. Si, "A method for suppressing SAR radio frequency interference based on deep learning in time-frequency domain," in *2023 International Conference on Image Processing, Computer Vision and Machine Learning (ICICML)*. IEEE, 2023, pp. 780–784.

- [4] J. Zhao, Y. Wang, G. Liao, X. Liu, K. Li, C. Yu, Y. Zhai, H. Xing, and X. Zhang, "Intelligent detection and segmentation of space-borne SAR radio frequency interference," *Remote Sensing*, vol. 15, no. 23, p. 5462, 2023.
- [5] N. Hendy, M. Czerkawski, N. Longép  , H. M. Fayek, and A. Al-Hourani, "Detection of radio frequency interference in raw sar data using u-net segmentation," *TechRxiv*, 2024. [Online]. Available: <https://www.techrxiv.org/users/706268/articles/1245246-detection-of-radio-frequency-interference-in-raw-sar-data-using-u-net-segmentation>
- [6] M. Czerkawski, C. Ilioudis, C. Clemente, C. Michie, I. Andonovic, and C. Tachtatzis, "Interference motion removal for doppler radar vital sign detection using variational encoder-decoder neural network," in *2021 IEEE Radar Conference (RadarConf21)*, 2021, pp. 1–6.
- [7] J. Shen, B. Han, Z. Pan, G. Li, Y. Hu, and C. Ding, "Learning time–frequency information with prior for SAR radio frequency interference suppression," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2022.
- [8] X. Li, J. Ran, H. Zhang, and S. Wei, "Mcsnet: A radio frequency interference suppression network for spaceborne SAR images via multi-dimensional feature transform," *Remote Sensing*, vol. 14, no. 24, p. 6337, 2022.
- [9] S. Wei, H. Zhang, X. Zeng, Z. Zhou, J. Shi, and X. Zhang, "CARNet: An effective method for SAR image interference suppression," *International Journal of Applied Earth Observation and Geoinformation*, vol. 114, p. 103019, 2022.
- [10] Y. Feng, B. Han, X. Wang, J. Shen, X. Guan, and H. Ding, "Self-supervised transformers for unsupervised SAR complex interference detection using canny edge detector," *Remote Sensing*, vol. 16, no. 2, p. 306, 2024.
- [11] Y. Huang, L. Zhang, J. Li, M. Tao, Z. Chen, and W. Hong, "Machine learning methods for SAR interference mitigation," in *Synthetic Aperture Radar (SAR) Data Applications*. Springer, 2022, pp. 113–146.
- [12] E. U. S. Programme, "Copernicus browser," 2022, accessed: 2024-01-09. [Online]. Available: <https://browser.dataspace.copernicus.eu>
- [13] N. Hendy, F. G. Kurnia, T. K. M. Bachmann, M. Martorella, R. J. Evans, M. Zink, H. M. Fayek, and A. Al-Hourani, "Semus-an open-source rf-level SAR emulator for interference modelling in spaceborne applications," *IEEE Transactions on Aerospace and Electronic Systems*, 2024.
- [14] A. Monti-Guarnieri, D. Giudici, and A. Recchia, "Identification of c-band radio frequency interferences from sentinel-1 data," *Remote Sensing*, vol. 9, no. 11, p. 1183, 2017.
- [15] M. Schandri, T. Kraus, A. Hourani, N. Hendy, F. Kurnia, and M. Bachmanna, "Analysis of radio frequency interferences in TerraSAR-X products," in *15th European Conference on Synthetic Aperture Radar. EUSAR*, 2024, pp. 6730–6733.
- [16] N. Franceschi, A. Recchia, R. Piantanida, G. Hajduch, P. Vincent, M. Pinheiro, N. Miranda, and C. Albinet, "Operational RFI mitigation approach in Sentinel-1 IPF," in *EUSAR 2022 - 14th European Conference on Synthetic Aperture Radar*, Leipzig, Germany, 2022, p. pp.
- [17] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A convnet for the 2020s," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11 976–11 986.