Karlsruhe Institute of Technology

Deutsches Zentrum für Luft- und Raumfahrt

Master Thesis

Analysis and Mitigation of Radio Frequency Interferences in X-band SAR Data

B.Eng. Caterina Simone Maria Klumpp

1st Examiner:	Prof. Dr.-Ing. Thomas Zwick
2nd Examiner:	Prof. Dr.-Ing. Marwan Younis
Supervisor:	Dr.-Ing. Markus Bachmann
	M.Sc. Thomas Kraus, M.Sc. Maximilian Schandri
Period:	11.12.2024 – 02.05.2025

Karlsruhe, 02.05.2025

Postal address: Institute of Radio Frequency Engineering and Electronics
Kaiserstraße 12
D-76131 Karlsruhe

Building: Engesserstraße 5, Geb. 30.10

Tel.: +49 (0) 721 608 4 2522
Secy.: +49 (0) 721 608 4 2523
Fax: +49 (0) 721 608 4 5027
Email: info@ihe.kit.edu
Web: www.ihe.kit.edu

Declaration

I hereby declare that I wrote my Master Thesis on my own and that I have followed the regulations relating to good scientific practice of the Karlsruhe Institute of Technology (KIT) in its latest form. I did not use any unacknowledged sources or means and I marked all references I used literally or by content.

Karlsruhe, 02.05.2025

B.Eng. Caterina Simone Maria Klumpp

Abstract

In recent years, terrestrial Radio Frequency Interference (RFI) from sources such as airports, harbors or air surveillance systems, as well as spaceborne RFIs from growing multisatellite constellations have increasingly impacted TerraSAR-X data.

This thesis investigates methods to detect and mitigate these interferences in raw data. Therefore, machine learning methods, including Isolation Forest (iForest), autoencoder, and Variational Autoencoder (VAE) architectures, are evaluated and compared regarding their detection performance. A clearly superior detection performance of the autoencoder for this particular case is found.

The performance of conventional Signal-to-Interference Ratio (SIR), Cell Averaging (CA)-Constant False Alarm Rate (CFAR) filtering and the autoencoder are further assessed and compared in terms of detection performance and mitigation effectiveness. As a reference a conventional SIR filter of $18\,\mathrm{dB}$ is determined for the scenes at hand and set as the benchmark. In case of Narrow Band Interferences (NBI) the CA-CFAR detects RFIs in additional $0.6\,\%$ of the image's samples and the autoencoder an additional $1.3\,\%$. In case of Wide Band Interferences (WBI) the CA-CFAR detects $7.93\,\%$ less than the SIR filter and the autoencoder an additional $2.7\,\%$. Therefore, the detection rate of the autoencoder not only outperforms the conventional methods but it is also more robust. However, blanking the isolated RFIs in each rangeline's Power Spectral Density (PSD) for mitigation shows that the improved detection rate of the autoencoder to the SIR benchmark visually has a neglectable impact.

Finally, a data set derived from gloablly distributed SAR acquisitions is analyzed based on occupied bandwidth, RFI power, and the number of RFIs in its PSD. This enables a global mapping of RFI occurrences in TerraSAR-X data from 2023 and 2024, along with an initial analysis of the nature of these RFIs and hence proves the power of the developed autoencoder.

Contents

List of Figures

List of Tables

Acronyms and symbols

Acronyms

2D	two-dimensional
ADC	Analog-to-Digital Converter
Adam	Adaptive learning rate optimization algorithm
AI	Artificial Intelligence
AUC	Area Under Curve
CA	Cell Averaging
CFAR	Constant False Alarm Rate
CM	Chirp Modulated
CNN	Convolutional Neural Network
CUT	Cell Under Test
CW	Continuous Wave
DC	Direct Current
DEM	Digital Elevation Model
DL	Deep Learning
DVB-T	Digital Video Broadcasting- Terrestrial
ESA	European Space Agency
FD	Frequency Domain
FDOA	Frequency Difference of Arrival
FFT	Fast Fourier Transformation

FN False Negatives

FP False Positives

FPR False Positive Ratio

FT Fourier Transformation

GMI Global Precipitation Measurement Microwave Imager

GNSS Global Navigation Satellite Systems

GPU Graphics Processing Unit

GRSS Geoscience and Remote Sensing Society

HPBW Half Power Beamwidth

I In-Phase

iForest Isolation Forest

iTree Isolation Tree

ITU International Telecommunication Union

KL Kullback Leibler

LDCR Lobe Differencing Correlating Radiometer

LEO Low Earth Orbit

LOS Line-of-Sight

MAE Maximum Absolute Error

MAPE Maximum Absolute Percentage Error

ML Machine Learning

MEO Medium Earth Orbit

MSE Mean Square Error

MUSIC MUltiple SIgnal Classification

NBI Narrow Band Interferences

NN Neural Networks

OFDM	Orthogonal Frequency Division Multiplexing
OS	Order Statistic
PCA	Principle Component Analysis
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Intervals
PSD	Power Spectral Density
PSLR	Peak-to-Side-Lobe Ratio
Q	Quadtrature
RAM	Random Access Memory
RCS	Radar Cross-Section
ReLU	Rectifier Linear Unit
RF	Radio Frequency
RFI	Radio Frequency Interference
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
RX	Receiving Antenna
SAR	Synthetic Aperture Radar
SIR	Signal-to-Interference Ratio
SINR	Signal-to-Interference-to-Noise Ratio
SLC	Single Look Complex Datatake
SM	Sinusoidally Modulated
SMAP	Soil Moisture Active/Passive
SMOS	Soil Moisture and Ocean Salinity
SNR	Signal-to-Noise Ratio
SVM	Support Vector Machine

TD	Time Domain
TDOA	Time Difference of Arrival
TN	True Negatives
TP	True Positives
TPR	True Positive Ratio
TTC	Telemetry, Tracking and Command
TX	Transmitting Antenna
UHF	Ultra High Frequency
VAE	Variational Autoencoder
WBI	Wide Band Interferences

Constants

c_0= $299\,792\,458\,\mathrm{m/s}$	Speed of light in vacuum
e= $2.718\,28$...	Euler's constant
π= $3.141\,59\ldots$	Pi
j= $\sqrt{-1}$	Imaginary unit

Latin symbols and variables

Lowercase letters

a	Amplitude (generic)
b	Detection scale factor CFAR Bias
b	Bias (Autoencoder, VAE)
c	Speed of light (used in context)
d_a	Real antenna length in azimuth
e	Attribute (decision trees)
f	Frequency
f_c	Center frequency
$f_\mathrm{i}(t)$	Instantaneous frequency
g	Gain
i	Index for summation (sample index)
k	Chirp rate (frequency modulation slope)
k_0	Chirp rate
l	Index for RFI component
m	Frame number (DVB-T)
n	Time/sample index
N	Length of time series (for RFI in signal)
N	Number of reference cells (CA-CFAR)
N	Dimension of latent space (VAE)
p	Split value (decision trees)
q	Attribute (decision trees)
r	Radial distance to target
r_0	Slant range (radial distance to target)
s_BP	Bandpass signal
t	Time
t	Number of trees (decision trees)
x	Signal sample
x	Point (decision tree)
x	Input (Autoencoder, VAE)
x_0	CUT CFAR
x_{iN}	Reference cells CFAR
y	Output (Autoencoder, VAE)
z	Latent dimension (Autoencoder, VAE)

Capital letters

A	Amplitude of RFI component
B	Scaling factor in CFAR thresholding
BW	Bandwidth
BW_{c}	Bandwidth of the chirp
D	Input Units (Autoencoder, VAE)
E	Energy of pulse
E	Estimated noise level
E	Average (decision trees)
G	Antenna gain
I	Interference (general RFI model)
K	Max subcarrier index
L	Total number of RFI components
L_{VAE}	VAE's Loss
L_{sa}	Synthetic aperture length
$L_{\mathrm{sa,unf}}$	Unfocused synthetic aperture length
M	Hidden Units (Autoencoder, VAE)
N	Length of time series (for RFI in signal)
N	Number of reference cells (CA-CFAR)
N	Dimension of latent space (VAE)
N	Count (generally)
P	Power
Q	Input signal (Autoencoder, VAE)
S	Received signal (true echo)
SIR_{az}	Signal-to-Interferer Ratio over the azimuth dimension
T	Threshold value for detection
T	Node for decision trees
T_{l}	Left daughter node (decision trees)
T_{r}	Right daughter node (decision trees)
W	Noise / Weight Matrix (Autoencoder, VAE)
X	Interfered signal (measured)
Y	Output signal (Autoencoder, VAE)

Greek symbols and variables

α	Attenuation constant (mentioned in context)
β	Modulation factor (SM RFI model)
δ	Resolution
ϵ	Random value of Gaussian distribution (VAE)
γ	Chirp rate (CM RFI model)
λ	Wavelength
μ	Mean
ϕ_f	Maximum phase error (in unfocused SAR)
ψ	Sub-sampling size
σ	Standard deviation
σ_{RCS}	Radar cross-section of the target
τ	Duration
θ_a	Beamwidth
$\varphi(t)$	Signal phase as a function of time
$\mathbb{R}\mathrm{e}$	Real part operator
j	Imaginary unit (also written as $\sqrt{-1}$)

Operators and mathematical symbols

$\equiv$	Equivalent
$\exp$	Exponential function
FFT	Fast Fourier Transform
$\frac{d}{dt}$	Time derivative
$\hat{\cdot}$	Predicted value
$\in$	Set membership
$\log$	Logarithm
$\odot$	Elementwise multiplication
$\mathbb{R}\mathrm{e}$	Real part operator
$\sin$	Sine function
$\sum$	Summation symbol
$\bar{\cdot}$	Arithmetic mean
$\geq$	Greater than or equal

General indices

BP	Bandpass
c	Center (in the case of frequency)
c	Chirp (in the case of bandwidth)
i	Instantaneous
P	Pulse
PSD,rl	Power Spectral Density in rangeline dimension
RCS	Radar Cross Section
RFI	Radio Frequency Interference
RX	Receiver
S	Signal
sa	Synthetic aperture
sa,unf	Unfocused synthetic aperture
SC	Subcarrier
SIR,az	Signal-to-Interferer Ratio in azimuth dimension
TX	Transmitter

1 Introduction

The use of spaceborne remote sensing using Synthetic Aperture Radar (SAR) and microwave radiometry techniques above optical sensing on satellites has increased strongly in the recent years and is increasing continuously. The advantage of those technologies is that they are independent of weather or the sun's light and therefore can be used for Earth observation at practically all times.

This real-time Earth observation data is crucial for the monitoring of natural disasters impact, e.g. earthquakes [Ell20], volcano eruptions, tsunamis [HOS+20]. Its availability is required and used by local authorities in order to coordinate disaster management and control [Air25]. Additionally, it contributes significantly to scientific fields in geosciences and the monitoring of climate change and natural resources. Ice shelves, glaciers and permafrost can be measured from space even in very remote and hard or not at all fully coverable areas such as Antarctica [KBC+14]. Furthermore, forest heights and types can be acquired in rain forest areas. Additionally, SAR technologies were used in order to acquire a global Digital Elevation Model (DEM) using SAR interferometry. Especially here those remote areas as Antarctica are impossible to measure terrestrially on this large scale with this precision [KFH+05]. SAR images are also used for agricultural or urban planning in general, such as infrastructure monitoring and security applications [Air25].

In the recent years the RFI occurence in X-band has been increasing. This is already known as a problem for passive spaceborne remote sensing particularly in the protected L-band [VDG+19]. But in recent years also increasing impact on active remote sensing in the protected X-band is rising. These interferers are estimated to be even more severely above the legally restricted limit. This is because instead of the passive radiometers measuring brightness temperatures emitted as reflections of cosmic radiation in general, SAR measures the backscatter of its own system's transmitted power. This means that the SAR signal's power levels are already higher than those of radiometers and the RFIs are still strong enough to dominate and contaminate the scene locally [SGS18]. This is even after further processing steps including matched filtering to the transmitted signal's pulse shape. However, lower frequency bands as the L-band also are generally easier accessible and more crowded which makes them more vulnerable to RFIs [dHHA19].

RFIs are mostly terrestrial and occur mainly around harbours, airports, military bases. However, there are also some RFIs caused by nearby satellites transmitting in X-band for Earth monitoring and flying in a similar orbit. This risk of RFIs caused by neighboring satellites is also increasing because of the current trend in the aerospace industry to rather build many smaller and cheaper satellites such as cube-sats. Therefore, the total number of satellites in the orbits increase. This increase is strongest in the Low Earth Orbit (LEO) rising from 1000 satellites in 2014 to 10 000 satellites in 2024 [Uni25]. This is already problematic itself. The TerraSAR-X and TanDEM-X satellites that are considered here also operate in the LEO and in X-band. However, the spectrum allocation for the X-band in LEO is additionally crucial for downlinking and disseminating space-collected data under any weather conditions which makes it even more problematic especially for the case at hand. This is caused by the fact that its downlink is far less expensive, complex and ultimately more adoptable than alternatives [Air25].

This increase in RFIs also in X-band SAR data makes it inevitable to analyze and search for RFI detection and mitigation techniques that can be reliably applied. Since the problem of RFI occurence and along the urge for documentation and RFI detection and mitigation techniques is already existing in other scientific fields for a long time, their approaches need to be considered first.

As an example RFI cleanliness is a long ongoing and consisting issue for radio astronomy and astrophysical observatories since very sensitive sensors are employed in order to detect very low power signals of events far away. However, this sensitivity also makes them more sensitive for RFIs which makes it unsurprising that unsupervised Machine Learning (ML) techniques as clustering are employed in order to detect RFI events from actual events [HCRDR19]. Another approach from radio astronomy employs edge-thresholding in real-time data to achieve RFI mitigation also in real time [BS19].

In Earth observation the issue of RFI detection and mitigation is investigated best and longest for passive remote sensing in radiometers. This evolution of the RFI scenario world wide, and its impact on the mission as well as lessons learned in RFI detection and geo-location and the problems of RFI contamination of the SMOS mission in L-band were discussed after 10 years in orbit already in 2019 [LDO+19]. Even though for example Soil Moisture Active/Passive (SMAP) has early step RFI filtering implemented already the detection is continuously improved in order to detect insufficiently detected hotspots using an iso-line mapping approach [SLVdM19]. SMAP's RFI detectors include pulse, cross-frequency, kurtosis, and polarimetric methods, as a function of the RFI sources' power and are regularly investigated for their performance. Also methods to examine remaining RFIs after detection and filtering are already established [BJMP17]. Furthermore, automatic RFI detection, clustering, identification and localization were researched in detail using long-term serial cross-polarization data source. For this use case an iterative density-based spatial clustering algorithm with noise and emission intensity characteristic that simultaneously considers the emission intensity and distribution density of the RFI source was suggested [WWF+19]. Overall, the three passive remote sensing satellites Aquarius, SMOS and SMAP in L-band can be compared and normally observe the same RFI distribution. They all have different established on-board detection and mitigation techniques implemented which might be the future for X-band SAR systems as well. While SMOS relies on a threshold detection algorithm, SMAP on a kurtosis detector and Aquarius on an outlier detector. Because of malfunction about Japan's distributed TV system and the aforementioned comparison, the outlier detector of Aquarius was however extended by an absolute detection threshold. This threshold is based on the maximum expected level of natural emissions [SdMLV16]. This and further, rather regular updates in the RFI detection and mitigation steps of radiometers are clearly documented and analyzed as for example in [dMSLV16] separating the steps in the active and passive instruments of Aquarius. Another example for the global importance and established integration of RFI detection and mitigation processing techniques in L-band is given in the Lobe Differencing Correlating Radiometer (LDCR) using a cross-frequency peak detection, time and frequency domain kurtosis, and complex coherence phase to detect RFIs in the radiometric data [VDG+19]. As a consequence of the increasing amount of RFIs for radiometers this effort in algorithms and hardware for RFI detection and mitigation is clear. However, due to the long development times of spaceborne systems the required onboard detection and blanking of RFIs was not implemented yet in former satellites in orbit but rather planned for and implemented in future systems like MetOp Second Generation and the Copernicus High Priority Mission Passive Imaging Microwave Radiometer satellites [KSSB19].

There is also research on detecting and classifying RFIs using ML algorithms for SMAP to obtain changes of RFI environment in L-band from 2015 to 2022 using a Convolutional Neural Network (CNN) [BBJ+22]. Beyond that research [PPQC24] is also suggesting ML techniques for the classification and characterization of real RFI events in passive remote sensing such as radiometry and Global Navigation Satellite Systems (GNSS)-reflectometry. Due to the amount of degrees of freedom this classification regarding identification and location of the interference source cannot be achieved adequately using classical signal processing techniques. For classification here CNNs and Recurrent Neural Network (RNN)s are employed and a future outlook is given on examining transfer learning and autoencoders for further improvements.

However, there are also RFI detection and mitigation techniques developed and in operation in SAR systems, particularly in Sentinel-1 in C-band already. For detection and mitigating whole RFI containing image parts, [LZL+22] is working on Level-2 data which is focussed imaging data. A method is presented that suggests first constructing an RFI-free background image from pre-processed time-series images from the past and then generating difference images from it with the image change detection method. The difference images entropy is calculated and compared to an adaptive threshold which replaces the conventional screening step. Subsequently, the preliminary RFI part is removed and the entropy calculation step of the new difference images is repeated. They yield the final screening and detection results and validate them on Sentinel-1 images. The aforementioned conventional screening step for Sentinel-1 was implemented since 2022 an interference detection module was added to the Sentinel-1 ground processor. This module includes pre-screening and detection functions which require the SAR system to receive additional noise signals. The detection is implemented using distribution statistics to locate the position of interference in the echoes. This is simple but comes with a low application accuracy due to the variety of interference types. In [FPR+21] however RFI is detected in Level-0 or raw data, which reduces the loss and impact of the mitigation. They use noise like rank echoes for global RFI monitoring. Those rank echoes are the first measures of each burst taken in Sentinel-1 in TOPSAR mode because the backscatter has not been received yet. RFI detection is implemented using the Fisher's Z for narrow-band high-power RFI and Kullback Leibler (KL) divergence for wide-band low-power RFI. Retrieved from the raw data's statistical properties this yields a global mapping of RFIs in

C-band.

Also Deep Learning (DL) based RFI detection is investigated for C-band RFIs. For example [IL19] investigates mimicking human visual inspection of processed images to detect RFIs in the measurement in instrumentation radar systems using a CNN. However, this is first of all also operating in C-band and secondly also implemented on processed imaging data which is mainly useful for labelling or for mitigation by mitigating or removing the whole contaminated part of the image which leads to enhanced loss of information.

This work focuses on the analysis of RFIs in X-band SAR data and additionally works on Level-0 meaning raw data. This aims for a high detection performance and a mitigation keeping the loss of information as low as possible.

RFIs are conventionally not only investigated on real RFI-contaminated SAR images. Especially, if they are still relatively rare as in X-band SAR data. Instead RFI simulators are built where the RFIs can be modelled and the effectiveness of detection, characterization and mitigation techniques can be compared on that particular modelled type of RFI. For example there is such a simulator for the Galileo uplink operations including its Ultra High Frequency (UHF) link for SAR, its S-band link for Telemetry, Tracking and Command (TTC), and its C- Band link for mission data. It includes RFI transmitters on ground and RFIs from satellites allowing sweeps of the interferer's power. The detection can be implemented as energy detection with a power detector in Time Domain (TD) or gaussian detector or a PSD detection with a space domain detector or a power detector in Frequency Domain (FD). The detected RFIs can be characterized by mean frequency, pulse width, occupied bandwidth, duty cycle and spectral kurtosis. For classification DL techniques such as RNN, CNN, Support Vector Machine (SVM) are used which require a ground truth labelled data set of the different RFI classes which is only feasible with simulated RFIs. Furthermore, the RFI can be localized using Time Difference of Arrival (TDOA), Frequency Difference of Arrival (FDOA) or the MUltiple SIgnal Classification (MUSIC) algorithm. It has to be noted that despite their usefulness for modelling the dependence on an RFIs power and for classifications real-world interferers differ significantly from models. Also, their variety is not represented nor do they meet most simulators' models' assumptions such as isotropic and continuously transmitting interferers [Oy22].

The target of this work is to analyze the detection, characterization and mitigation of RFIs in real X-band SAR data of the TerraSAR-X mission. In order to do so ML techniques and architectures are to be selected, implemented and compared regarding their performance for the detection of RFIs.

In Chapter 2 firstly the theoretical foundations are built. They focus on the SAR principle and on RFIs, conventional RFI detection, introducing DL methods and methods for RFI mitigation. Next, in Chapter 3 the three differently used data sets for different types of analysis are introduced. Subsequently, the DL architectures' selection is described along with their pre-processing steps, design choices and the implementation in Chapter 4. With the final comparison of their performance the best DL architecture is chosen for further investigations and analysis on the X-band SAR data sets.

The chosen DL architecture is tested on conditions representative to the real-world RFIs and compared with conventional methods to benchmark in Chapter 5. Further classification and isolation of the RFIs in the PSDs yields a global classified RFI mapping regarding total occupied bandwidth, RFI power and the distribution of separate RFI counts in the PSD. Finally, having isolated the RFIs spectrally the comparison with conventional methods in detection rate can be validated on processed images using frequency blanking as mitigation technique in Chapter 6.

2 Theoretical Fundamentals

This chapter is treating all of the theoretical fundamentals required to understand the following results. In the beginning the basic function of radar systems and imaging radars will be discussed before introducing the SAR for spaceborne radar remote sensing. Additionally, the considered ML concepts and basic functions and applications relevant for the further work will be illustrated.

2.1 Imaging Radar and Spaceborne Radar Remote Sensing

There are several options to obtain images for Earth observation from space. The most basic concept is to employ commercial camera principles in order to obtain optical images [LWZ19]. This method is the most cost-effective one in terms of achieving favourable results. However, it is subject to certain constraints in terms of time, namely the requirement for good weather conditions and daylight hours.

Another passive sensor based principle is spaceborne microwave radiometry, which is measuring cosmic background and thermal radiation backscattered from Earth. Here a receiving antenna is employed as a sensor. The measurement is wide-band operated and therefore containing atmospheric windows. Therefore, it is weather independent and radiation from weather effects in the atmosphere as clouds or rain can be detected just as well as lower frequent radiation from underlying structures. However, the resolution of a radiometer image is about 1000-10 000 times lower than for a SAR image. Therefore, the information of the structures, vegetation and soil can be joined but radiometers alone cannot be used for a precise mapping. In spaceborne radar remote sensing the active sensor principle is utilised. Therefore, the antenna as a sensor is replaced by a radar system operating the antenna in transmit and receive mode. This is how imaging can take place independent from the sunlight since the measured backscattered power is the one transmitted by the system. Weather dependent effects are avoided and the penetration depth is set by selecting the operating wavelength of the system as a design choice. That way the system can be optimized to measure atmospheric effects as clouds or Earth surface effects, such as even penetrating into vegetation or soil [CM91].

2.1.1 Basic Radar Concepts

Radar sensors are being used for distance measurement in the first place. In case of radar remote sensing pulsed and coherent radars have to be used in order to form the synthetic aperture. The following chapter will provide a brief introduction of the aforementioned radar types.

Pulsed Radar

Operating a pulsed radar system a high frequency pulse that is short in TD is transmitted and reflected by the scene. The radar echo is detected as the received power in the frequency band used for transmitting. This can be described mathematically by the radar equation Eq. 2.1 where G is the antenna gain of the Transmitting Antenna (TX) and the Receiving Antenna (RX) and λ the according wavelength of the chosen frequency of the sent pulse. r is the radial distance of the target to the radar system which functions as a transmitter (TX) as well as as receiver (RX). The target is being characterized by its Radar Cross-Section (RCS) σ_{RCS}. Knowing these parameters the received pulse power P_{RX} is calculated from the transmitted pulse power P_{TX} as

$$P_{\text{RX}} = \frac{G_{\text{TX}} G_{\text{RX}} \lambda^2 {\sigma_{\text{RCS}}}^2}{(4\pi)^3 r_{\text{TX}}^2 r_{\text{RX}}} P_{\text{TX}}. \tag{2.1}$$

A pulsed radar system keeps sending pulses equidistant in time domain with a fixed Pulse Repetition Frequency (PRF). This PRF determines the receiving window, in which a received pulse can be allocated unambiguously to the transmitted pulse. A pulse received in the next receiving window will be allocated to the next transmitted pulse and therefore is interpreted to be a closer target.

The pulsed radar systems resolution is determined by the pulses duration τ_{P} or the according bandwidth $BW = 1/\tau_{\text{P}}$. It can be derived from the minimum geometric distance of two targets that can be differentiated for their echoes not to overlap at the receiver [KH10]. This is expressed as

$$\delta_r = \frac{c_0}{2BW}. \tag{2.2}$$

This signifies that the radar system's ability to separate neighboring targets is limited by the bandwidth of the transmitted signal's pulse. Physically close targets cause the received pulses to overlap in time domain. An increased bandwidth or shorter pulse duration therefore leads to an increased resolution, reducing the probability of the pulses to overlap.

However, a signal's energy is determined integrating the signal's power over time. Applying the reduction of the pulse length therefore directly translates into a reduction of the signal's energy E as

$$E = P_{\mathrm{TX}}\tau_{\mathrm{P}}. \tag{2.3}$$

This trade-off between resolution and received Signal-to-Noise Ratio (SNR) requires a high pulse power which makes the concept itself solely not applicable for most cases.

Frequency Modulated Coherent Pulsed Radar

Combining the pulsed radar with a frequency modulated pulse shape allows separating physically close targets along with long pulse duration by their different instantaneous frequencies [Mor00]. The transmitted signal is shaped as a linear frequency ramp which is called a chirp and can be described mathematically as

$$s_{\mathrm{BP}} = \sin(2\pi(f_{\mathrm{c}} + kt)t), \text{ with } t \in \left[-\frac{\tau_{\mathrm{P}}}{2}, \frac{\tau_{\mathrm{P}}}{2}\right]. \tag{2.4}$$

The signal is centered around the center frequency f_{c} and the chirp rate k is calculated as

$$k = \frac{BW_{\mathrm{c}}}{\tau_{\mathrm{P}}}. \tag{2.5}$$

In order measure the range the received signal's phase shift to the transmitted signal's phase needs to be determined. Therefore, the instantaneous frequency $f_{\mathrm{i}}(t)$ is derived by differentiating the signals phase $\phi(t)$ after time as

$$f_{\mathrm{i}}(t) = \frac{1}{2\pi}\frac{\mathrm{d}\varphi(t)}{\mathrm{d}t} = \frac{1}{2\pi}\frac{\mathrm{d}}{\mathrm{d}t}\left[2\pi\left(f_{\mathrm{c}} + kt\right)t\right] = f_{\mathrm{c}} + \frac{BW_{\mathrm{c}}}{\tau_{\mathrm{P}}}t. \tag{2.6}$$

Here it is important to consider that the radar system's operating bandwidth needs to operate at double the chirp frequency in order to fulfill the nyquist criterion. This particularly concerning the sampling frequency of the Analog-to-Digital Converter (ADC).

In order to maximize the received SNR a matched filter is applied. The matched filter is a digital filter convolving the transmitted signal with its complex conjugate in TD or multiplying it in the FD [KH10].

Radar systems are mainly characterized by their impulse response. It is described by the radar systems response to a normalized point target. From this the geometric resolution can be determined from the 3-dB-bandwidth of its main lobe. The sidelobe suppression can be evaluated from the Peak-to-Side-Lobe Ratio (PSLR). The basic reference function is a $sinc$-function which without any weighting function employs a PSLR of $13\,\mathrm{dB}$. This is often improved by using a hamming-weighting-function to the antenna pattern which leads to a PSLR of $43\,\mathrm{dB}$, so that there is more contrast in the resulting image. However, this comes to the cost of a wider main lobe and therefore a slightly decreased resolution [Mor00].

2.1.2 Synthetic Aperture Radar

A frequency modulated pulsed radar can be enhanced to a two-dimensional (2D) imaging radar system. In order to do so the system is mounted onto a moving platform e.g. a satellite in orbit and a synthetic aperture is being processed from the received backscattered signals.

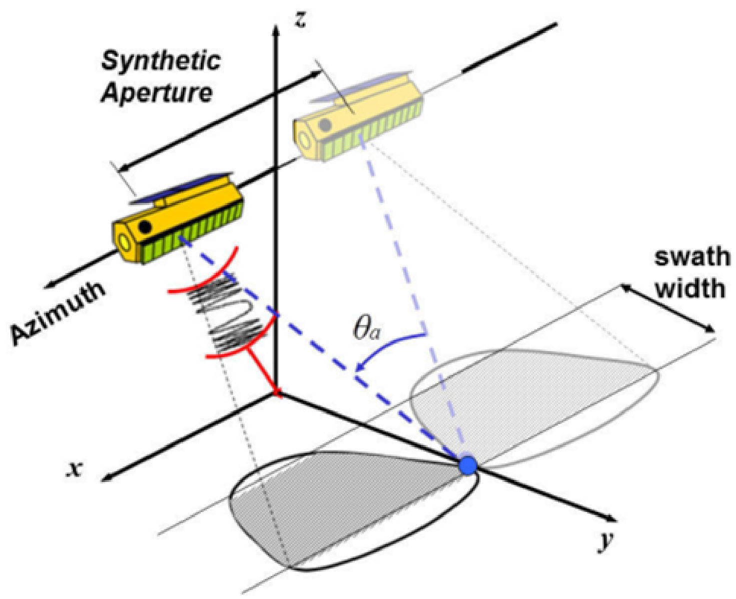

Figure 2.1: Geometric principles of a SAR scene including the azimuth and range direction, the swath width and footprint of the SAR antenna's beam, the transmitted chirp and the SAR system's synthetic aperture θ_{a}.

SAR systems can operate in a monostatic, bistatic or multistatic manner, depending on the amount of receivers and their location relative to the transmitter. Fig. (2.1) shows the geometry of a SAR scene. On the Earth the radar system's footprint can be seen that is enlighted by the radar system's beam when imaging the target. The target is a hypothetical point of interest on the Earth's surface. However, the real SAR system measures the beam footprint as an area instead of the hypothetical point target.

The flight direction is called azimuth and the perpendicular Line-of-Sight (LOS) beam direction to the target is called slant range or range direction. The point on the Earth's surface right below the satellite is called the nadir. With this point another range measure can be taken which is the according ground range along the Earth surface to each slant range [CW05].

In the following work when referring to the beamwidth the half-power beamwidth is meant. The beamwidth in elevation or range direction causes the beam footprint on the Earth's surface to spread an area from the nearest point in near-range to the furthest in far-range. This spread corresponds to the swath width on the ground [CW05].

Since the radar system is mounted on a platform moving in azimuth direction SAR systems utilize the SAR Doppler frequency. The Doppler effect describes that the received signal's frequency increases when the distance to the target decreases and on the other hand decreases with an increasing distance to the target.

In order to make use of this effect it is necessary to operate the radar system in a pulsed manner and coherently with a constant PRF. If a coherent operation is not ensured, meaning the start time and phase of the pulses is not accurately controlled, the information of the phase center and therefore the option to process the synthetic aperture is lost. Therefore the Doppler SAR frequency works as a function of the carrier frequency that the baseband signal is upconverted to before it is transmitted more than as a function of the baseband signal itself [CW05].

A real antenna's angular resolution θ_a in azimuth direction is described by the beamwidth of its main lobe. This resolution depends on the wavelength λ and the length of the antenna d_a only

$$\theta_\mathrm{a} = \frac{\lambda}{d_\mathrm{a}}. \tag{2.7}$$

However, this is undesirable, since the resolution can only be improved by decreasing the wavelength and therefore increasing the frequency and increasing the antenna length. Both of those parameters however are limited, since the higher frequency goes along with a higher sampling frequency. This sampling frequency needs to be double the chirp's bandwidth $BW_($c according to the nyquist criterion. Also the physical antenna length mounted on a spaceborne system is limited by practical reasons.

Forming the synthetic aperture radar by creating synthetic antenna elements in between the imaging time the resolution can be adapted to θ_sa. Instead of depending on the antennas dimension depends on the length of the synthetic aperture L_sa as

$$\theta_\mathrm{sa} = \frac{\lambda}{2L_\mathrm{sa}}. \tag{2.8}$$

The factor two appears because the of the time shift in transmitting from each synthetic element that is considered "two-way". Since in a real antenna or array all elements transmit at the same time the gradient in phase only appears on the backscattered path of the signal. By forming the synthetic aperture however the path of the transmitted signal also experiences a phase shift. Nevertheless, the phase shift is equal to the one on the way back since it is caused by the transmitted signal's path length which is equal to the backscattered signal's path.

The length of the synthetic aperture is defined by L_{sa} in Eq. 2.9:

$$L_{\mathrm{sa}} = \theta_{\mathrm{a}} r_0 = \frac{\lambda}{d_{\mathrm{a}}} r_0. \tag{2.9}$$

Inserting L_{sa} from Eq. 2.9 into the product of θ_{sa} from Eq. 2.8 and r_0 yields the synthetic aperture's resolution in range direction

$$\delta_{\mathrm{sa}} = \theta_{\mathrm{sa}} r_0 = \frac{d_{\mathrm{a}}}{2}. \tag{2.10}$$

This means that the resolution in azimuth direction becomes equal to the antenna resolution itself and also invariant to range [Mor00].

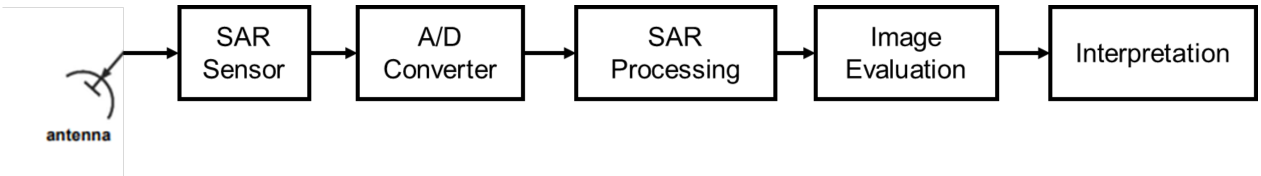

Figure 2.2: In the principle SAR data flow the power transmitted and received by the antenna front-end at different times passes the RF electronics and amplification stages in the SAR sensor before being A/D sampled and demodulated into the digital domain. In the following the digital raw data is range- and azimuth compressed by applying matched filters to the transmitted chirp and Doppler and the raw data is focused to be imaging data. The final image can be evaluated and interpreted further.

Summarizing the SAR processing to one step in a data flow chart the signal's processing chain is depicted in the block diagram in Fig. 2.2. In the SAR sensor which is containing the Radio Frequency (RF) electronics the signal is generated and pulse modulated and by a circulator switched to the transmitter circuit to be amplified and transmitted via the antenna front-end. Switching back to operate the antenna in the receiving mode during the echo windows in the Pulse Repetition Intervals (PRI) in between transmitting the pulses, the antenna receives the echos that are amplified and demodulated by the SAR sensor's RF circuit. Therefore the received pulse is bandpassfiltered, amplified and downconverted to baseband. The demodulation takes place in the sampling step, sampling the In-Phase (I) and Quadtrature (Q) components [CM91]. Furthermore the signal is sampled according to the nyquist theorem in order to preserve all information which means a minimum sampling rate of double the signal's maximum frequency component. In order to fulfill the first nyquist criterion the sampling frequency in range direction, which is also defined as the fast time direction because it is determined by the pulse duration, would be expected to be minimally bigger than double the chirp's bandwidth. However, it can be made use of the I/Q modulation of the signal, which allows to reduce the sampling frequency by a factor of two since the signal is split in two channels, in its real and imaginary part. Therefore in elevation a sampling frequency only higher than the chirp's bandwidth is sufficient. The same effect is exploited when sampling in azimuth direction so the sampling frequency is set by the PRF. The azimuth direction is also named slow time because it is determined by the illumination time going along with the synthetic aperture length. Therefore it needs to be larger than the significant Doppler bandwidth. Since the azimuth spectrum rolls of more slowly than the range spectrum conventionally an oversampling factor in between 1.1 and 1.4 is chosen in order to avoid aliasing and ambiguities [CW05].

Having the signal downconverted to baseband and sampled in digital domain, it enters the SAR processing chain. As already discribed the pulse shape that is chosen to be transmitted for the SAR signal is a chirp, usually chosen to be an upscaling chirp. Therefore, the received signal in range direction is expected to be shaped as the transmitted signal. In order to implement a matched filter which optimizes the SNR the received signal is convolved with a reference signal. This reference signal is the complex conjugate of the received signal's shape and therefore the complex conjugate of the transmitted signal. This processing step is called range compression.

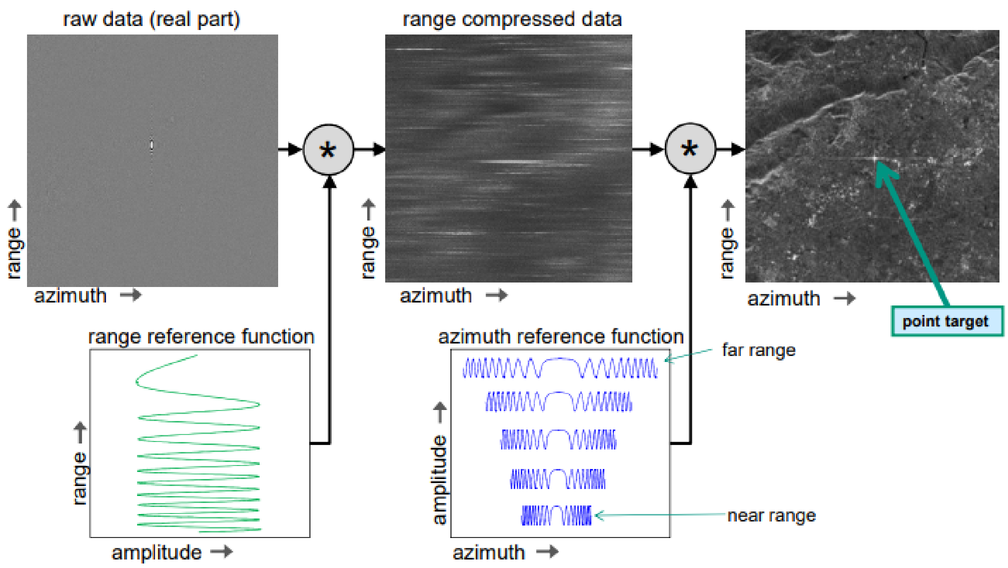

Figure 2.3: SAR processing chain with the 2D raw SAR data experiencing range compression and azimuth compression in order to focus the scene and particularly a marked point target. This involves the convolution with the according reference function and considers the range variance in case of the azimuth compression.

In azimuth direction the Doppler effect which is caused by the moving platform has a similar modulation effect taking place. Therefore a similar azimuth compression step is implemented convolving the range compressed data with an azimuth reference function following the concept of the matched filter. This can be seen in Fig. 2.3. However, it is important to notice that in contrast to the range effects the azimuth effects are range variant. Therefore, the azimuth reference function needs to be range variant too [Mor00].

In order to enable near-realtime processing the concept of unfocused processing evolved. This concept means reducing the length of the synthetic aperture. This is done in a way so that the far-field condition from Eq. 2.11 for coherent processing is possible without focussing:

$$r_0 \geq \frac{2d_{\mathrm{a}}^2}{\lambda}\,. \tag{2.11}$$

In unfocused processing the phase is considered but not corrected. The received signal is allowed to have phase errors up to 90°. The higher the phase errors the higher are the sidelobes in the impulse response. Skipping the correction of the phase error the reference function is rectangularly shaped in the signal's bandwidth. Therefore the convolutional integral simplifies to integrating the azimuth signal over the illumination time that is defined by the length of the synthetic aperture. With the maximum allowed phase error being set to 90° [Mor00] the unfocused length translates to

$$L_{\text{sa,unf}} = \sqrt{\frac{2\lambda r_0 \phi_f}{\pi}} \equiv \sqrt{\lambda r_0}. \tag{2.12}$$

For the resolution this goes along with deterioration and range variant effects [Mor00] as can be seen

$$\delta_{\text{sa,unf}} = \frac{L_{\text{sa,unf}}}{2} \equiv \frac{\sqrt{\lambda r_0}}{2}. \tag{2.13}$$

2.1.3 Radio Frequency Interference

Interfering signals emitted from other radio frequency sources are occuring in the received radar data as RFI. It is common that SAR systems operating at lower frequencies as in P- or L-band or also operating with a larger bandwidth are particularly sensitive to RFIs [TSHW19].

As well for active remote sensing with SAR but even more so for passive remote sensing with radiometers in the lower frequency bands (< $15\,\mathrm{GHz}$) the occurence and impact of RFIs in and therefore on the imaging data is very severe. In the case of SAR the presence of RFI in an acquisition which is a so called "data take" may corrupt information including amplitude, frequency, time delay, polarization, Doppler shift and phase [TSHW19]. Therefore, the system's frequency bands are conventionally protected by regulations proposed on international level by the International Telecommunication Union (ITU) [GCW$^+$15].

According to the frequency allocation rules [Uni09] imposed by the recommendation ITU-R RS.577-7, the allowed frequency bands and bandwidths for space-borne SAR sensors is listed in Tab. 2.1.

The Geoscience and Remote Sensing Society (GRSS) is providing a global map showing RFI

Table 2.1: The current spectrum and bandwidth regulatory rules for each space-borne active sensor set by the ITU including band designation, allowed frequency interval and permitted bandwidth.

Band designation	Allowed frequency interval	Permitted bandwidth
P band	432 to 438 MHz	6 MHz
L band	1215 to 1300 MHz	85 MHz
S band	3100 to 3300 MHz	200 MHz
C band	5250 to 5570 MHz	320 MHz
X band	8550 to 8650 MHz	100 MHz
	9300 to 9900 MHz	300 MHz

detections for different instruments and frequency bands as shown in Fig. 2.4. It can be noted that in well-developed and densely populated areas as Europe or southern Asia the RFI occurence is a lot more crowded. Since this work is handling X-band SAR data particularly the X-band RFIs ranging from 8 GHz to 12 GHz are of interest. However, all only data of NASA's Global Precipitation Measurement Microwave Imager (GMI) radiometer operating at 10.6 GHz is entered in the underlying data base.

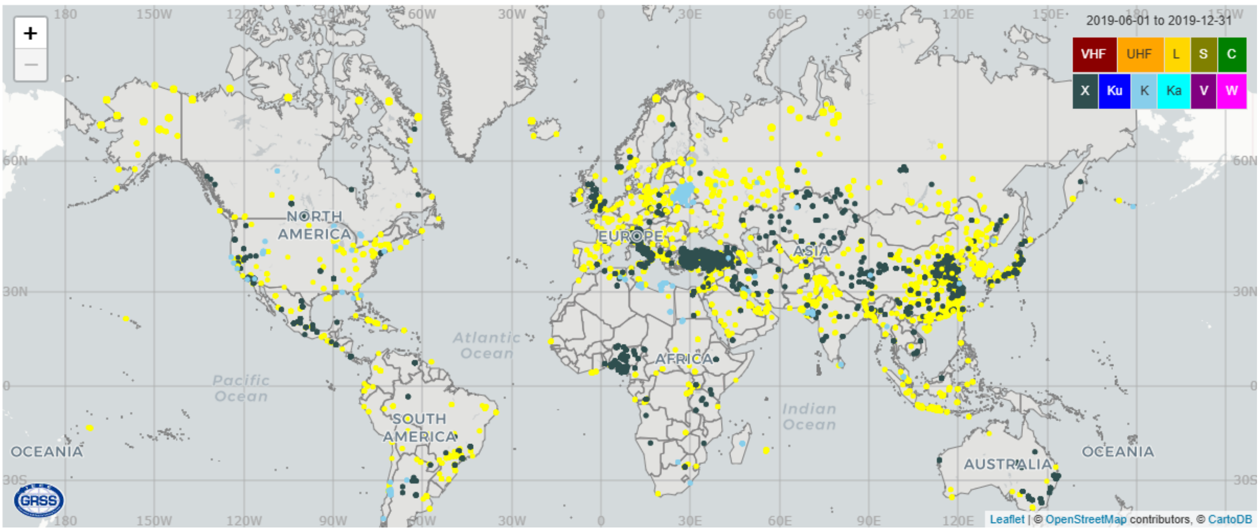

Figure 2.4: Globally detected RFIs by scientific missions are reported and registered by the IEEE GRSS. This scientific database is made accessible by the visualizations of the IEEE GRSS RFI observations display system [GRS19].

The SAR data this work is based on is yielded from the TerraSAR-X and TanDEM-X mission of the German Aerospace Center. Those two satellites fly in close formation on a polar orbit around Earth and operate in X-band at a center frequency of 9.65 GHz with bandwidths ranging from 100 to 300 MHz. They can be operated in a monostatic or bistatic mode depending on the case of application. For generating SAR images or in bistatic mode using interferometry in order to create

DEMs for example. They can also be operated in different imaging modes and polarizations. Using different polarizations has the advantage to get more information about the surface structure and materials, and imaging modes may enhance the monitored area or resolution. The resolution and monitored area depends highly on the operating mode. FOr example for large areas are covered by the ScanSAR mode with a lower resolution while for high resolution but smaller areas the Spotlight mode is used. However, filtering the X-band RFIs in the map above to the operating bandwidth of the TerraSAR-X Mission which is ranged $9.5\,\mathrm{GHz}$ to $9.8\,\mathrm{GHz}$ the number of documented RFIs in the GRSS Map reduces to zero, since there were no RFIs reported from satellites operating in that frequency range yet.

In order to avoid conflict of remote sensing systems to collide to the maximum extent, the radio transmission in each frequency band is categorized into primary and secondary services by the ITU. Depending on the frequency band SAR can therefore be a primary basis in one band but only be a secondary basis in another one and therefore be denied from imaging large areas. This is the case in P-band where it is also only allocated by a permitted bandwidth of $6\,\mathrm{MHz}$. Additionally, it is the cause of restricting the European Space Agency (ESA)'s biomass mission from large areas of Northern America and Europe because of possible interferences with wind profilers and space object tracking radars [TSHW19].

Most RFI sources are originated in commercial or industrial radio devices over land. These include for example radiolocation radars, wind profilers, telecommunication devices and amateur radios [TSHW19]. In X-band they can be seen most prominently observed in areas of airports, harbours and military stations. However, there are also few cases of RFI originated in other space-borne satellites, such as broadcasting signals from the GNSS constellations, communication satellites or other active remote sensing systems. These spaceborne RFI sources can affect the SAR system in two ways. One is hitting the SAR system with a direct jamming signal within an antenna side-lobe. This kind of interference power was analyzed and found out to be non-negligible but tolerable due to the antenna pattern. The other one is terrain scattered interference that occurs when two SAR satellites share a nearby illuminated area. This interference gets potentially large since the main lobes could couple and hit specular targets on specular surfaces as lakes. Then the received echo of a target would not be seen as a point-like source anymore. The combination of the both SAR system's transmitted signal's and backscatter would overlay and cause a remodulation across the illuminated area. This scenario was not actually found to have occured yet. However, it is gaining in probability for future missions with large swath widths, more frequent revisits and orbit crossings [TSHW19].

The received echo $S(n)$ is just like the RFI $I(n)$ and the noise $W(n)$ a time series and can therefore be modeled of length N with elements n. Therefore the interfered Signal $X(n)$ can be mathematically modeled as the summation:

$$X(n) = S(n) + W(n) + I(n), \text{ with } 1 \leq n \leq N \,. \tag{2.14}$$

Low power incoherent RFI could be "mitigated" during the matched filtering in SAR processing. However, this is more a smearing than a controlled filtering. Therefore, high power RFIs remain in the focused image as visible artifacts and noise.

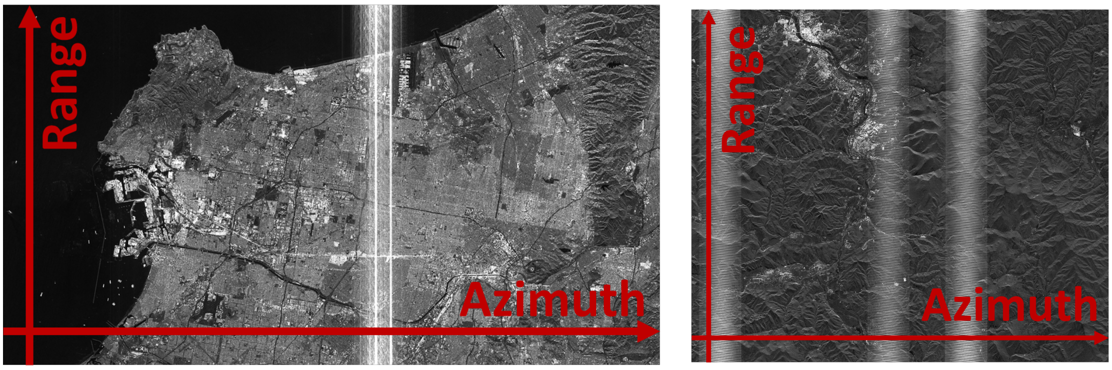

(a) LA, USA (b) Erdaojiang, China

Figure 2.5: Exemplary focussed 2D SAR images. They contain different RFI patterns and are aligned with the azimuth axis lying in the x-axis and the range axis lying on the y-axis. This convention is kept for the following work.

The SAR data is 2D- imaging data with the azimuth and range axis. One rangeline is therefore defined to describe all range samples per azimuth sample. These visible artifacts are in azimuth direction normally limited by the antenna pattern but in range direction blurred over the whole rangeline.

If there are a multitude of low power RFI sources in bistatic data the injected noise causes the loss of the phase information and introduces a phase error that may make it unapplicable for SAR interferometry [TSHW19].

NBIs take a narrow bandwidth in relation to the SAR's chirped signal bandwidth and are continuous in time relatively to the illumination time. This type of RFI usually occurs in land-mobile and amateur radio systems and can be modeled mathematically as a sum of complex sinusoidal tones A_l, f_l and θ_l being amplitude, center frequency and phase for the l-th RFI component [TSHW19]:

$$I_{NBI}(n) = \sum_{l=1}^{L} A_l(n) \exp(j(2\pi f_l n + \theta_l)). \tag{2.15}$$

Pulsed WBIs have a varying PRF and a wider bandwidth. They are common in ground-based radiolocation radars. They can be differentiated into Chirp Modulated (CM) (Eq. 2.16) and Sinusoidally Modulated (SM) (Eq. 2.17) WBI, which normally each do not describe real RFIs perfectly but real RFIs can rather be described as a combination of both of them:

$$I_{\text{CM}}(n) = \sum_{l=1}^{L} A_l(n) \exp(j(2\pi f_l n + \pi\gamma_l n^2)), \tag{2.16}$$

$$I_{\text{SM}}(n) = \sum_{l=1}^{L} A_l(n) \exp(j\beta_l sin(2\pi f_l n + \phi_l), \tag{2.17}$$

with γ_l, β_l and ϕ_l being the chirp rate, modulation factor and initial phase of the l-th component.

Broadband Continuous Wave (CW) RFI signals are modulated and therefore occupy a relatively broad bandwidth. Also they are continuous in time with respect to the integration time. They often occur in broadband communication systems or coded signals as the GNSS and have noise like characteristics. This could be for example caused by a Digital Video Broadcasting- Terrestrial (DVB-T) Signal as in Eq. 2.18 that is organized in frames and sub-carriers using Orthogonal Frequency Division Multiplexing (OFDM) symbols. With m, l and k_{SC} being the frame number, symbol number and subcarrier number. The information s is modulated on the subcarriers and in order to extract the actual physical signal only the real part is retrieved [TSHW19]. This type of RFI is the hardest to detect and mitigate.

$$I_{\text{DVB-T}}(n) = Re\left\{\exp(j2\pi f_c n) \sum_{m=0}^{\infty} \sum_{l=1}^{l_{max}} \sum_{k_{\text{SC}}=1}^{K_{\text{max}}} s\varphi(n)\right\}, \tag{2.18}$$

with:

$$\varphi(n) = \begin{cases} \exp(\frac{j2\pi k(n-T_g-lT_S-l_{\text{max}}mT_{\text{s}})}{T_u}) & \text{if } (l+l_{\text{max}}m)T_S \leq n \leq (l+l_{\text{max}}m+1)T_S), \\ 0 & \text{else.} \end{cases} \tag{2.19}$$

The signal is transmitted in frames. Each fram duration T_S consists of data part T_u and a guard interval with the duration T_g. Each frame consists of $l_{\text{max}} = 68$ OFDM symbols.

Fig. 2.6 shows a comparison the normalized magnitudes of those three RFI types logarithmic in exemplary range spectra.

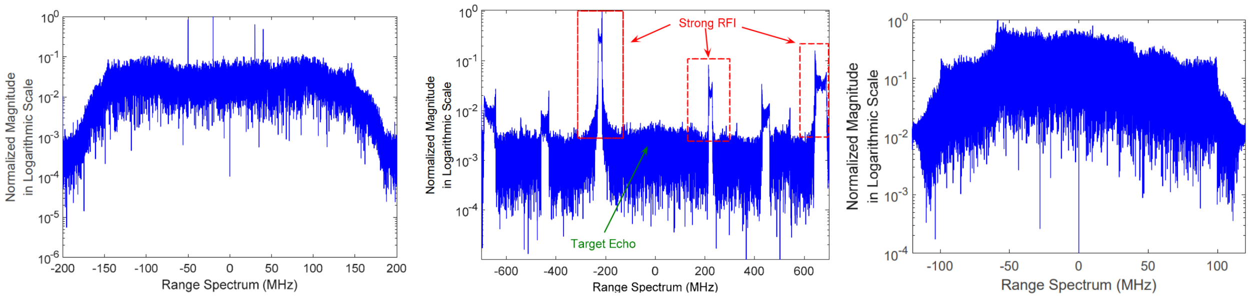

(a) Narrow-band (b) Pulsed Wide-band (c) Broadband

Figure 2.6: Differing exemplary RFIs from [TSHW19] in their normalized range spectra in comparison to each other.

2.2 Conventional RFI Detection

RFI detection methods can be separated according to their point of action in the data flow. This split results into the categories of detecting RFIs in level-0 data which effectively means unfocused raw data or level-1 data meaning the focused imaging data [LZL$^+$22].

RFI detection in the level-0 data can be implemented in TD, FD or the time-frequency-Domain. The easiest implementation is using a power threshold and comparing the received signal in TD. This is adequate for short and strong pulses, however weaker longer RFIs remain undetected.

A more adaptible way of detection is to transform the received signal to the FD using a Fast Fourier Transformation (FFT) and comparing it to a threshold or perform detection algorithms in the FD.

However, most conventional is the detection and mitigation with spectrograms in time-frequency-domain. Thereby the signals dimensional extension in time and frequency is taken into account. So continuous wave and pulsed RFIs can be detected simultaneously [PQPC20].

A conventional approach to separate targets from background noise is the CFAR. It was originally suggested by Finn and Johnson in [Fin68] and is since then state-of-the-art in many variations depending on the scenery.

The simplest case of a CFAR is the CA-CFAR. The CA-CFAR is working according to Fig. 2.7. There is one Cell Under Test (CUT) x_0 considered at a time that is surrounded by guard cells on each side. The noise level estimate E is calculated by averaging the reference cells x_{iN} on both sides where N is the length of the CFAR filter meaning the number of cells considered as

$$E = \frac{1}{N} \sum_{i=1}^{N} x_i. \tag{2.20}$$

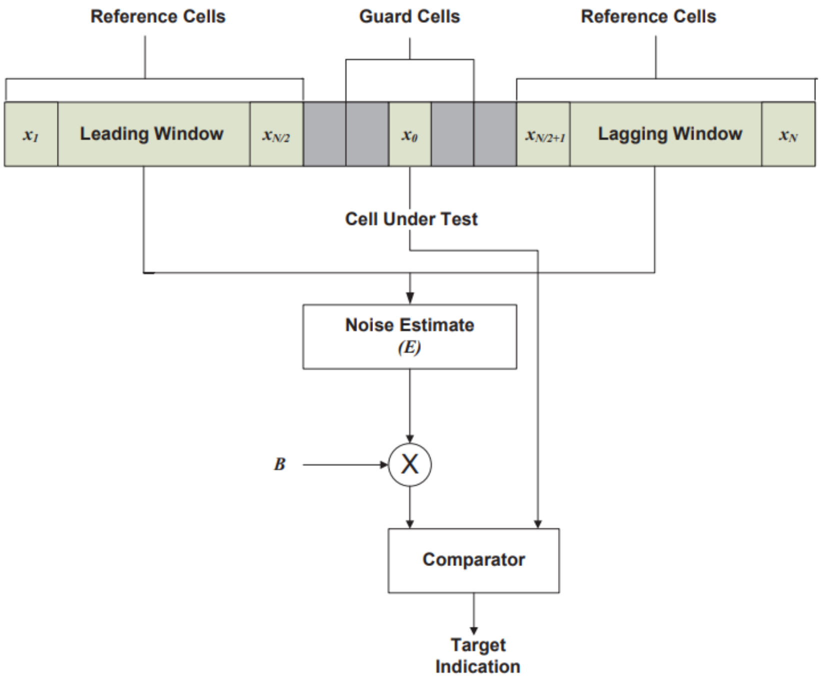

Figure 2.7: Basic function principle of a CA-CFAR from [JYB16] including the CUT x_0, the guard cells and the reference cells on but sides. It also depicts the resulting noise estimate E and the multiplication with the detection scale factor B leading to the final target indication.

Scaling this noise estimate by a predetermined detection scale factor B the threshold T can be calculated by $T = EB$. With this threshold T there is a binary decision for every CUT x_0 if there is a target or RFI detected or not [JYB16].

Calculating the sliding window over the backscatter and deciding on pixel basis means that CA-CFAR processing yields large computational effort and low speeds. In order to accelerate the computation, the fast CA-CFAR is a CA-CFAR enhanced by a convolution of the CFAR window and the signal's power in the raw data or image's intensity in the imaging data instead of pixel by pixel performance [YLL$^+$24].

CA-CFAR is a very effective detector in case of stationary and homogeneous interferers. However, in case of multiple targets or RFIs some can be masked by the biggest target or strongest RFI [JYB16]. Also in the case of WBIs the averaging results in layering around the target which deteriorates the detection probability.

A CFAR algorithm that is supposed to deal with that problem of separating multiple close targets is the Order Statistic (OS)-CFAR. In OS-CFAR the reference windows samples are sorted in ascending order and there is no guard interval. The k-th element of the sorted list is called the k-th order statistic, with the $n/2$-th order statistic being the data's median. False alarms can be reduced in case k-th order statistic is chosen close to the maximum, however that goes along with a loss of detection probability [JYB16].

2.3 Deep Learning Methods

In a conventional engineering design flow a problem is typically defined, analyzed and standardized by experts in order to implement a mathematical model that can be solved analytically or numerically. However, sometimes this design flow is too costly or inefficient, or the problem is too complex to be studied in its full generality. In this case ML and DL models have evolved. Instead of relying on domain knowledge ML relies on large data sets labelled according to ground truth of the task and trains and optimizes general-purpose learning machines accordingly to fulfill the desired task [S$^+$18]. The terminology DL evolves from the principle that the ability to learn complicated concepts is built out of multiple simpler ones. Depicting this structure as a graph there are many individual simple layers on top of each other and therefore showing a deep structure [GBC16].

Finding unexpected, rare patterns in data is defined as anomaly detection. In order to do so with ML the detector is trained with normal data only. This is highly advantageous since the anomaly's nature is sparse which often causes the problem that even with a large data set available the amount of anomalous data samples is insufficient for a supervised classification training. However, in anomaly detection scenarios usually a lot of normal data samples are available which enables the use of normal data for training in semi-supervised detectors. Even though more powerful methods for fault detection such as knowledge-based methods and supervised learning-based methods do exist, they are mostly not feasible in real scenarios because they require detailed a priori knowledge of potential faults and perturbations [GAFC23]. Therefore the conventional model-based approach with statistical methods, classification-based methods, and clustering-based methods for anomaly detection is to construct a profile of normal instances, to afterwards identify anomalies as instances that do not conform to the normal profile [LTZ09].

The field of DL is broad and the methods and architectures are numerous. However, they can be categorized most basic into supervised and unsupervised learning. In supervised learning a training data set with according target variables or so-called labels is given. The goal therefore is to find an algorithm to predict the target variable or label for a new data point on the exterior of the so far observed domain. The most prominent yet basic example for this type of learning is linear regression.

In contrary to that in unsupervised learning training is taking place on unlabelled data sets. Apart from that major difference however, unsupervised learning is more loosely defined. Given a data set the goal is to learn some useful information on the characteristics of interests for the specific application at hand from the data distribution given. The learning process typically uses hidden or latent variables. Example uses are density estimation for a direct approximation to the data set's distribution, dimensionality reduction and representation or the generation of new samples [S$^+$18].

There are various architectures of ML used for anomaly detection. So are e.g. instance based methods used frequently because they are simple, efficient and intuitive computing an outlier score from the distance of a new sample point relative to the data points stored in the memory. Prominent examples for this category are the k-nearest neighbor algorithm such as the local outlier factor [LTZ09].

While ensemble architectures can be as well used in combination with ML methods in order to increase performance or sensitivity they can also be used for anomaly detection on their own. The most well-known example for that is the iForest constructed by an ensemble of Isolation Tree (iTree)s for direct isolation of outliers [LTZ09]. This approach will be discussed in the following work.

Alternatively, explicit generalization models retrieve the normal behavior in a fixed size model, e.g. linear regression models, boltzmann machines or Neural Networks (NN)s while requiring low memory and offering various possible extensions of the basic methods. As part of explicit generalization models most prominently NNs that Artificial Intelligence (AI) falls into can be named. In anomaly detection most NN architectures are based on autoencoders, VAEs or time-series-networks, where the two prior ones will be explained in the following work in detail [GAFC23].

2.3.1 Isolation Forest

The iForest is a model-based architecture that explicitly isolates anomalies rather than profiling normal instances. It exploits the anomalies' attributes of being a minority consisting of fewer instances and having clearly differentiable attribute-values from normal instances. Because of those attributes anomalies in a decision tree are isolated closer to the roots whereas normal points are isolated deeper down at the end of a tree as is depicted in Fig. 2.8.

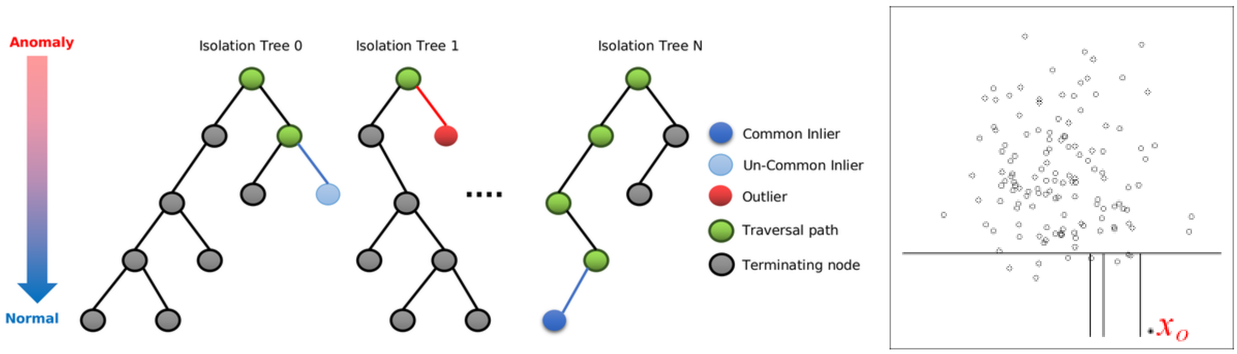

(a) Principle [DHS22]

(b) Isolating x_0 [LTZ09]

Figure 2.8: Basic function principle of the iForest network architecture showing its ensemble structure. Furthermore, the implementation of outlier isolation is shown.

The iForest is an ensemble consisting of several iTrees. Each iTree consists of a certain amount of nodes. Each node T can be either an external node with no child because the prior split already met the splitting condition. Or it is an internal node with one test and exactly two daughter nodes T_l and T_r. In an iTree every test consists of an attribute q and a split value p. The test is therefore splitting the data set at T along $q<p$ into the subsets of the data set at T_l and T_r. Implementing those splits an iTree is set up by recursively dividing the data set by randomly choosing an attribute q and a split value p until either a height limit is reached, the subset is containing one sample only or all of the subset's elements have the same value. If each node in the tree has exactly zero or two daughter nodes the iTree is a proper binary tree. Since all anomalies are expected to be isolated to external nodes of number n the memory requirement is bounded and only grows linearly with number of isolated anomalies [LTZ09].

In order to detect anomalies a ranking that reflects the degree of anomaly is required. One way to deal with this is to sort the data samples along path length and anomaly score and to pick the tops of the list as anomalies. The path length $h(x)$ of a point x is measured by the number of edges x traverses in an iTree from the root node until reaching an external node [LTZ09].

Deriving an anomaly score is challenging since while the maximum possible height of iTree grows in the order of n, the average height grows in the order of $log(n)$. However, the anomaly score is computed as in Eq. 2.21 from the average path length $E(h(x))$ of a collection of isolation trees and $c(n)$ which describes the average of path length $h(x)$ for a given external nodes of number n. This is borrowed from describing an unsuccessful search in binary search trees [LTZ09]:

$$s(x,n) = 2^{-\frac{E(h(x)}{c(n)}}. \tag{2.21}$$

Along the anomaly score s the sample can be ranked as an anomaly or not. If instances return s very close to 1, they are definitely ranked as anomalies. Instances with s much smaller than 0.5, are quite safe to be regarded as normal instances, while if all instances return $s \approx 0.5$, the entire sample does not really have any distinct anomaly.

Multiple iTrees for a given data set stuck together in an ensemble form an iForest. An iForest is a powerful method for anomaly detection and classifies based on the average path lengths, while its only variables are number of trees t and subsampling size ψ. Empirical studies in [LTZ09] have shown that for various applications choosing $\psi = 256$ generally provides enough details to perform anomaly detection with desired behaviour and convergence for $t \leq 100$. Anomaly detection using the iForest is implemented in a two-stage process. The first stage is the training stage which builds iTrees using sub-samples of the training set. The result is exemplary iTrees with decision nodes with particular attributes q and a split values p. The second stage is the testing stage which passes the test instances through iTrees to obtain an anomaly score for each instance. As described prior the iTrees are functioning in the recursive splitting manner building a partitional model. The limit in tree height is set by roughly the average tree height from training since only shorter than average path lengths are expected to be anomalies [LTZ09].

2.3.2 Autoencoder

Autoencoders are neural networks consisting of the same number of output units as input units. They are trained to generate an output y that is close to the input x. A trained autoencoder feedforwards each input to a unique representation $z(x)$ in the latent dimension within the neural network. Overall the autoencoder therefore consists of two parts, the encoder mapping x to the hidden representation $z(x)$ and the decoder mapping $z(x)$ onto the output $y(z)$. In order to find non-trivial solutions and learn significant features, some constraint is required so that the network does not simply copy the input values to the outputs. A conventional approach is to restrict the latent dimension of z relatively to the input x.

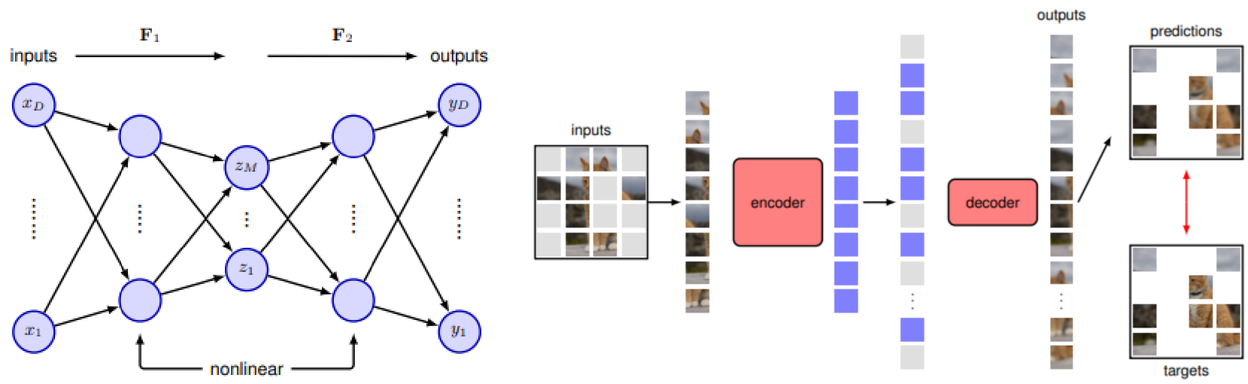

(a) Architecture (b) Example for image

Figure 2.9: Basic function principle of the autoencoder after [BB23]. The non-linear hidden units in the deep, layered encoder and decoder structure are depicted mapping the input to the output layer via the latent dimension. Additionally, an example implementation for the case of an image is given.

This compression is normally taking place over various layers and the architectures can therefore be described as $deepautoencoders$. Fig. 2.9 shows a basic multilayer network with D inputs and outputs mapped onto M hidden units and a dimensionality compression meaning $M < D$. The target variables for training are the input itself since the goal is to reconstruct the input values x from the latent dimension data $z(x)$. This is called an autoassociative mapping. The reconstruction error from input to reconstructions is described by an error function. Minimizing this error function the network parameters w are determined [BB23]. The encoded data $z(x)$ in the latent dimension can be described by the encoding layer's activation function f, the weight matrix W and the bias b as in [STG21]:

$$z(x) = f(Wx + b). \tag{2.22}$$

Reconstructing the data in the decoder layers is described symmetrically by decoding layer's activation function f', its weight matrix W' and its bias b' as

$$y(z) = f'(W'h + b'). \tag{2.23}$$

For linear activation functions the error function has a unique global minimum. At this minimum the network performs a projection onto the M-dimensional subspace that is spanned by the first M principal components of the data. This special case of a linear autoencoder is simple and very prominently used as the Principle Component Analysis (PCA) in order to choose the most significant attributes in complex data.

Including nonlinear layers in the network however, the input and output units, such as potentially even the latent layer units are still linear. However, changing the encoding and decoding layers to use sigmoidal nonlinear activation functions makes this mapping very general and not restricting linearity. This improves generalization abilities. However it also requires computationally intensive nonlinear optimization techniques that come with increased complexity of and introduce the risk of finding a sub-optimal local minimum of the error function [BB23].

2.3.3 Variational Autoencoder

Similar to the standard autoencoder, VAEs consist of an encoder, latent space and a decoder. They also aim to minimize the reconstruction error between the input data and its reconstructions. However, while a standard autoencoder encodes each input vector as a single vector in latent space, VAEs encode the input as a Gaussian distribution in the latent space [STG21].

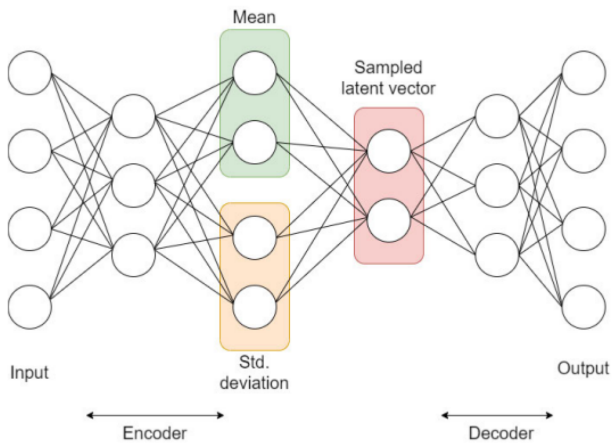

Figure 2.10: Basic function principle of the VAE after [STG21] keeping the autoencoders encoder and decoder structure. The gaussian distribution in latent space is represented by its mean μ and its standard deviation σ Additionally, the resampling trick with random sampling from this distribution in training is depicted. This is supposed to increase the model's generalization abilities and therefore also its robustness.

This principle is depicted in Fig. 2.10. The ML's backpropagation principle cannot be applied to a random sampling process when training the model. Therefore a 'reparameterization trick' is applied, where a random value is sampled from a unit Gaussian distribution [KW19].

The latent vector z is described by μ, σ, ϵ being the mean, standard deviation and a random value of the gaussian distribution. The $\odot$ is the elementwise multiplication [STG21] as

$$z = \mu + \sigma \odot \epsilon. \tag{2.24}$$

The VAE loss L_{VAE} function consists of a reconstruction and a regularization term. The regularization term is expressed as the KL divergence between the learned distribution and the true prior distribution. It is computed by μ, σ, x,y which are the mean, standard deviation, input and reconstruction as

$$L_{\text{VAE}}(x, y, \mu, \sigma) = L(x, y) \sum_{i=1}^{N} \sigma_i^2 \mu_{i=1}^2 - \log(\sigma_i) - 1. \tag{2.25}$$

$L(x, y)$ is the reconstruction term and $\sum_{i=1}^{N} \sigma_i^2 \mu_{i=1}^2 - \log(\sigma_i) - 1$ is the regularization term and N describes the dimension of the latent space [STG21].

2.3.4 Evaluation Metrics

In order to quantify and compare an architecture's performance evaluation metrics are required. One of the most commonly used is the accuracy. The accuracy is computed from the True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN) and is a first measure for the ability in anomaly detection

$$accuracy = \frac{TP + TN}{TP + TN + FP + FN}. \tag{2.26}$$

However, it has to be noted that the accuracy alone is not sufficient for particularly imbalanced problems such as anomaly detection. Since the data load for the normal data (negatives) is much higher than for the anomalous data (positives) a high accuracy score does not automatically mean a high anomaly detection rate. Instead further metrics have to be taken into account [GAFC23].

Another commonly used evaluation metric is the precision. The precision is used to describe the percentage of results that are relevant and is calculated as

$$precision = \frac{TP}{TP + FP}. \tag{2.27}$$

Consequently, the recall is the percentage of the relevant results that were classified correctly by the algorithm that is observed. It is defined as

$$recall = \frac{TP}{TP + TN}. \tag{2.28}$$

Ultimately the F1-Score calculates the balance between precision and recall from the two metrics above in order to design the algorithm adequately. It yields a value between zero and one with higher values indicating a balanced performance [STG21] as

$$F1 = 2\frac{precision recall}{precision + recall}. \tag{2.29}$$

Another two metrics particularly useful for anomaly detection is the True Positive Ratio (TPR) and False Positive Ratio (FPR) as computed in Eq. 2.30 and Eq. 2.31. Their importance resumes from their function in the Receiver Operating Characteristic (ROC)-curve. The ROC-curve optimizes the trade-off in between TPR and FPR, meaning the maximum correct detection rate of anomalies along with the minimum false detection rate of normal data samples [GAFC23]. They are defined as

$$TPR = \frac{TP}{TP + FN}, \tag{2.30}$$

and:

$$FPR = \frac{FP}{FP + TN}. \tag{2.31}$$

With the ROC-curve the threshold separating the normal and anomalous detected samples is swept and which thus optimizes the trade-off for the samples at hand. The resulting performance is conventionally depicted in a confusion matrix separating the normally and anomalously labelled samples according to their predicted labels. Such an exemplary confusion matrix is depicted in Fig. 2.11. These measures will be relied on in order to comprehensively compare the network architectures in the following work.

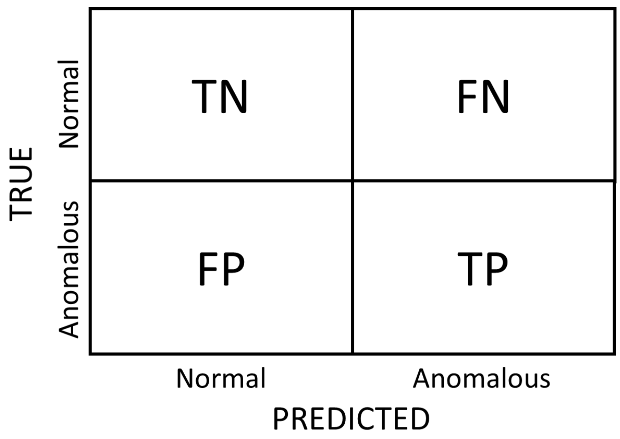

Figure 2.11: Exemplary confusion matrix showing the structure used for the later on performance measure. It separates truely normal or anomalous samples and visualizes the amount of them that are predicted to be normal or anomalous by the investigated model.

2.4 Mitigation

Whilst most research investigates the RFI detection and mitigation in the focussed SAR imaging data such as CFAR [YLL+24] and Notch Filtering in [ZLX+24], it also offers benefits to implement the detection and mitigation on the raw data directly. This is because the realtime processing abilities for rapid onboard processing of the SAR imaging data is limited. This requirement offers opportunities for the investigations of RFI detection and mitigation on the raw data as it is driven for example with a filter bank approach in [MACC24] and a two-stage mitigation with pulse-by-pulse difference detection in [WHG+25].

As in conventional detection RFIs in SAR data can also be mitigated in the TD, FD or time-frequency domain.

In TD Mitigation is taking place by techniques pulse blanking and amplitude domain processing. However, it is very limited to the well-functioning for short and strong power pulses.

In FD Mitigation is conventionally implemented by notch filtering or frequency blanking. Frequency blanking means locating the RFIs in the raw data spectrum and setting a power threshold to limit the RFIs' bandwidth sideways. Using frequency blanking the detected bandwidths in the rangelines are blanked out and set to zero, therefor the RFIs' power and but also the signal for these bandwidths is lost [PQPC20]. The most intuitive approach is to use a notch filter. The concept is simple yet very effective, which is why it is widely used for spaceborne onboard mitigation. Most notch filters consider a particular threshold and zero notch to suppress the RFI with the notch as a contrary shaped filter. There are also novel approaches trying to introduce adaptability to the RFIs' energy [ZLX+24]. However, the notches advantage depends on the actual RFIs' resemblance with the assumed RFI shape. Notch filters assume narrow-band RFIs which is the RFI type occuring most frequently.

In time-frequency fomain detection and mitigation go hand in hand conventionally using spectrograms and a frequency banks for mitigation [MACC24].

Adaptive filters are more complex and can be implemented in time, frequency, spatial, polarization, or by multi-domain joint analysis. The most popular one is the recursive least-mean-squares algorithm that can achieve a good compromise for convergence speed, stability, complexity and adaptability at the same time. The filter coefficients are calculated by minimizing the mean square prediction error. Another advantage is the computational effort and speed, since wide-band as well as narrow-band RFIs are mitigated within one filter's pass. This is enabling the mitigation technique for onboard processing which is another reason for its frequent usage [TSHW19].

3 Data Set

In order to detect, analyze and investigate RFIs in SAR X-band data first of all a data set needs to be built. Therefore, the scenes were chosen based on their quicklook images yielded from the TerraSAR-X mission in monostatic operation mode. For this analysis three types of data sets were built. The first one consists of selected imaging scenes without RFI and scenes with strong RFI that are distributed on a global level. The second is a subset of this data set in order to establish a validated ground truth for it. The third data set consists of noise pulses from all TerraSAR-X acquisitions performed in 2023 and 2024. They are used to expand the analysis to a global coverage.

3.1 Scene Selection

First of all quicklook images around the world containing and not containg RFI are chosen and separated according to whether their quicklook images were visually showing signs of RFIs or not. This sorting yielded 99 RFI-free scenes of which two examples can be seen in Fig. 3.1.

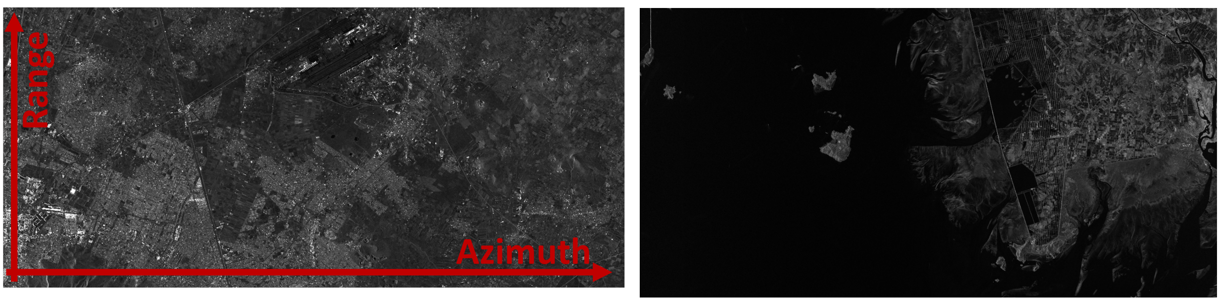

(a) Mexico City, Mexico

(b) On-Dong, North Korea

Figure 3.1: Three exemplary quicklook scenes of the data set without visible RFI artifacts. This data set consists of a total of 99 scenes scenes that are selected on a visual inspection basis.

Additionally, it yielded 28 scenes containing strongly visible RFIs. 16 of those scenes are located at Los Angeles. As a result of a previous study most of the scenes are found to contain interferences close to an airport's radar system at different times [SKAH$^+$24]. The variety in those scenes is led back to the radar system differing in operating frequency and mode over time even during one satellite overflight. Therefore, even for the same stationary RF interferer as a source the behaviour, interference with the SAR system and therefore resulting RFI pattern in the focussed image can vary a lot. However, there are various radar systems serving as RFI sources at, e.g., airports, harbours or military bases around the world. Since already the same one differs a lot in operation and RFI pattern, different RFI sources have an even larger variance. This can be seen in Fig. 3.2.

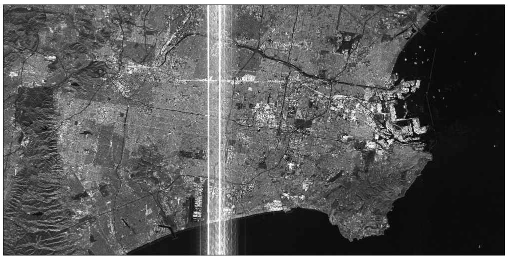

(a) Los Angeles, USA

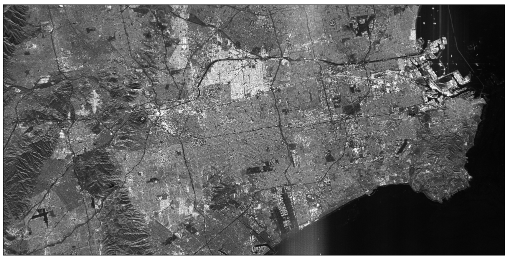

(b) Los Angeles, USA

(c) Liadong Bay, China

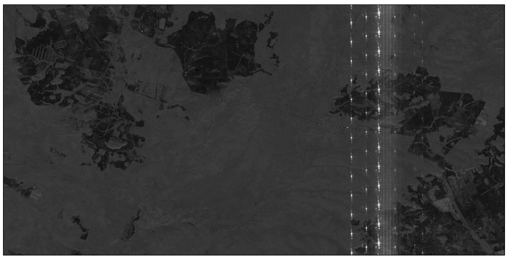

(d) Australia, Queensland

Figure 3.2: Four exemplary quicklook scenes of the data set with strongly visible RFI artifacts. This data set consists of a total of 99 scenes scenes that are selected on a visual inspection basis.

3.2 Pre-Study on Rangeline Level

What can also be seen in these examples in Fig. 3.2 is that not the whole scene is affected by the RFI. The RFI is smearing and destroying the information for the whole rangeline. Because of the azimuth compression it should be smeared in azimuth direction as well. However, due to the antenna pattern the dominant artifacts are very compressed to few rangelines in azimuth direction. Depending on the interference the antenna pattern and sidelobes of the SAR antenna can be noted, but the overlaying RFIs Doppler together with the SAR Doppler seem to compress the RFIs impact in SAR processing to few rangelines only. These RFIs in the processed images can be visually classified as RFI or RFI-free by experts easily but with high manual effort. Therefore, replacing this manual form of RFI detection by ML algorithms is already being investigated using CNNs [IL19]. CNNs divide images into areas of pixels and scan those pixels in order to detect significant areas for classification. It is known that for each affected rangeline the whole processed rangeline will be smeared. However, for scenes containing RFIs most of the image's rangelines are not affected by the RFI. Since mitigation can only be implemented in the raw data anyways also the detection is chosen to be implemented on raw data. Another advantage of choosing the raw data domain is using on-board detection in large bandwidth systems to choose other frequency bands for the data take as an immediate action. Since RFIs are strongly visible in the rangelines PSDs already, the sampled raw data $x[n]$ is converted into the Fourier domain using an FFT and the absolute value is taken. That way the PSD $S_x[k]$ by the FFT mapped to frequency bins k is computed. It is normalized over the total amount of N samples per rangeline and divided by the frequency bin width or spectral resolution Δf

$$S_x[k] = \frac{1}{f_\mathrm{s}}|\mathrm{FFT}(x[n])|^2 = \frac{1}{N\Delta f}\left|\sum_{n=0}^{N-1} x[n]e^{-j\frac{2\pi}{N}kn}\right|^2, \tag{3.1}$$

with the frequency bin width Δf described by the sampling frequency f_s and the total length of samples N:

$$\Delta f = \frac{f_\mathrm{s}}{N}. \tag{3.2}$$

The PSD $S_x[k]$ describes how the power of a signal or time series is distributed across different frequencies. It shows which frequencies contain the most energy, helping to analyze patterns like noise, periodic signals or RFI. Therefore, those rangeline PSDs are the foundation of the data set.

3.2.1 Imaging Scenes

A basic characteristic with which RFIs can be detected or classified in the PSD is by the SIR. A RFI signal can be characterized by its peak power

$$P_\mathrm{RFI} = \max|\mathrm{FFT}(x[n])|^2, \tag{3.3}$$

However, in order to interpret its impact it needs to be put into relation to the signal's power

$$P_\mathrm{S} = \mathrm{mean}|\mathrm{FFT}(x[n])|^2. \tag{3.4}$$

The SIR of a signal can be described as

$$SIR = \frac{P_\mathrm{RFI}}{P_\mathrm{S}}. \tag{3.5}$$

It can not only be used to quantify an interferer's strength in relation to the signal's strength in general. Another use is to classify the RFI-contaminated signals in terms of significance and notability in focused SAR. This allows to determine a measurable threshold underneath of which a potential RFI might be neglectable or is not be separable at all from the backscattered signals.

This filtering along a SIR threshold is implemented as a pre-detection of strong interferers. It is important in order to implement a filtering of only rangelines that are secured to contain RFIs of quantifiable strength to test and compare the detection and mitigation method's performance. Doing so the first challenge is to define a suitable threshold. This is achieved by calculating the SIR for each rangeline over the whole RFI scenes in azimuth direction and analyzing the according mean SIR $\mu_{\text{SIR,az}}$ and $\sigma_{\text{SIR,az}}$. Adding those two variables yields a first estimate of an upper threshold for the scene at hand as

$$T_{\text{SIR, az}} = \mu_{\text{SIR, az}} + \sigma_{\text{SIR, az}}. \tag{3.6}$$

In order to picture this threshold in an anomalous rangeline the computed $T_{\text{SIR, az}}$ is added to the mean PSD $\mu_{\text{PSD,rl}}$ as

$$T_{\text{PSD,rl}} = \mu_{\text{PSD,rl}} + T_{\text{SIR, az}}. \tag{3.7}$$

These thresholds are depicted exemplary in Fig. 3.3.

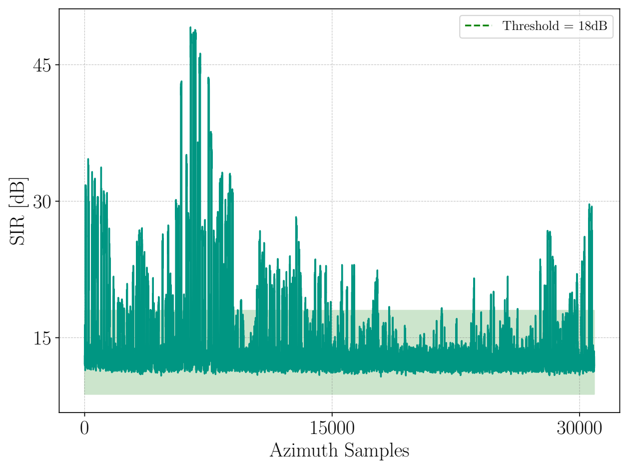

(a) SIR over azimuth

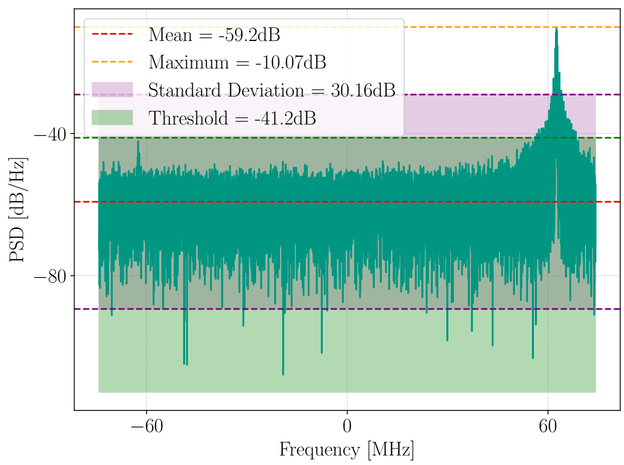

(b) PSD over rangeline

Figure 3.3: Filtering all rangelines of an RFI-contaminated scene by a SIR threshold of $18\,\text{dB}$. The threshold $T_{\text{az,rl}}$ is retrieved from all RFI-scene's mean SIR $\mu_{\text{SIR,az}}$ and its standard deviation $\sigma_{\text{SIR,az}}$. More deeply, the threshold $T_{\text{PSD,rl}}$ is depicted in an according PSD rangeline spectrum together with the PSD-specific maximum and mean value that correspond to the RFI's power P_{RFI} and the signal's power P_{S}.

Having this thresholding technique run over all 28 RFI scenes yielded $T_{\text{SIR, az}}$ in between $16\,\text{dB}$ to $18\,\text{dB}$, while having it run over all 99 no RFI scenes yielded $T_{\text{SIR, az}}$ with a maximum of $13\,\text{dB}$. It was validated empirically that a conservative threshold of $T_{\text{SIR, az}} = 18\,\text{dB}$ indeed yielded only RFIs with a notable SIR while lower thresholds included RFI-free rangelines as well.

Having this filtering also check the visually sorted scenes showed three presumably RFI-free scenes to contain RFI as well. They are depicted in Fig. 3.5. Comparing the quicklook images with the SIR over azimuth with the set threshold depicted in Fig. 3.4 not even weak RFI patterns can be found in the quicklook according azimuth positions in the quicklook images.

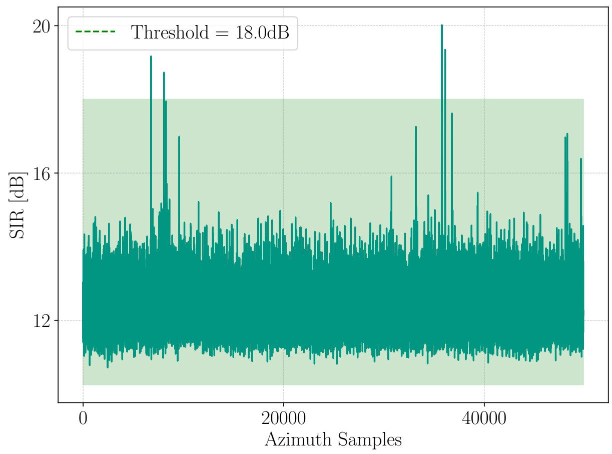

(a) Taqtaq Iraq

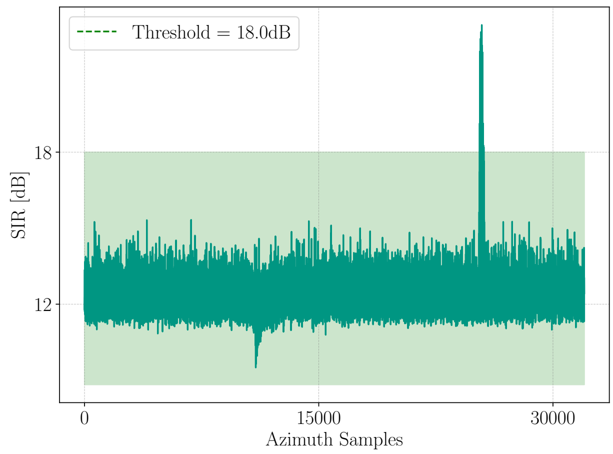

(b) Xuchang China

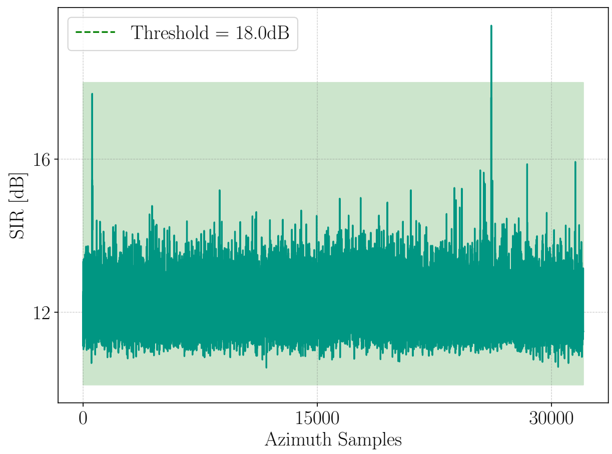

(c) Abdali Kuwait

Figure 3.4: Three RFIs found that are visible in the SIR filtered profiles but that are not visible in the according processed quicklook images shown in Fig. 3.5

However, the existence of interferers in the detected PSD rangelines is validated as can be seen in Fig. 3.6. Therefore, the three RFI-affected RFI-free files are sorted to the RFI scenes. Therefore the scene ratio switched to 96 RFI-free scenes to 31 secured RFI scenes.

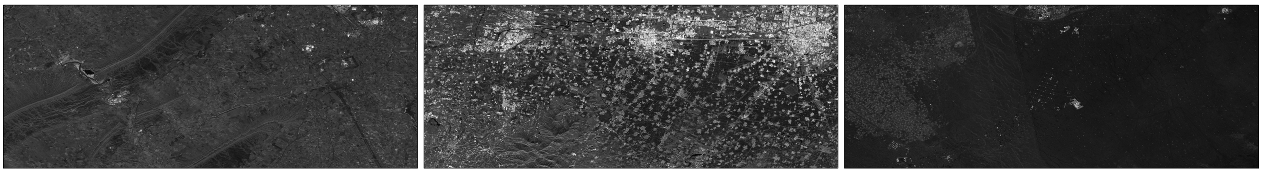

(a) Taqtaq, Iraq (b) Xuchang, China (c) Abdali, Kuwait

Figure 3.5: RFIs invisible in the quicklook images but visible in the according SIR profiles.

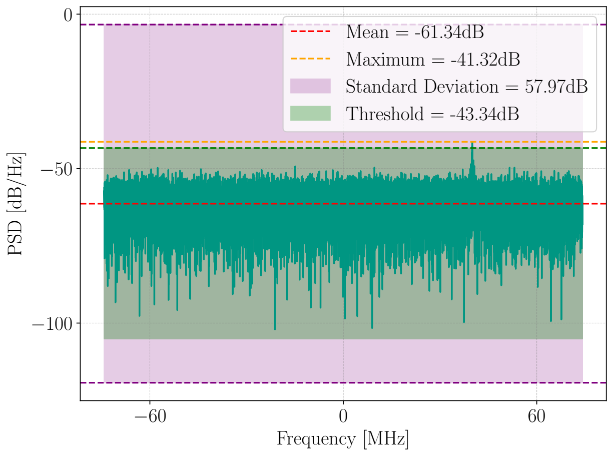

(a) Taqtaq, Iraq

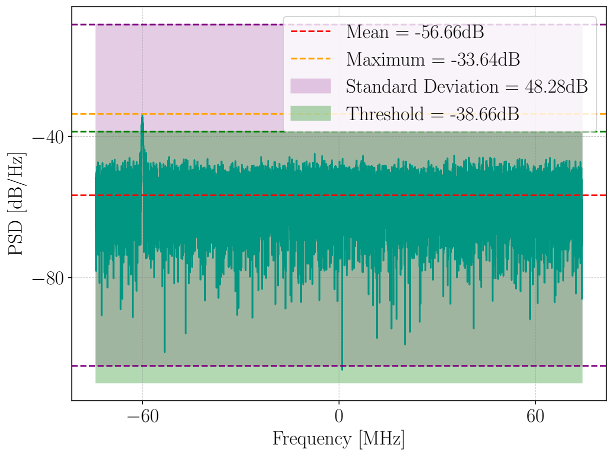

(b) Xuchang, China

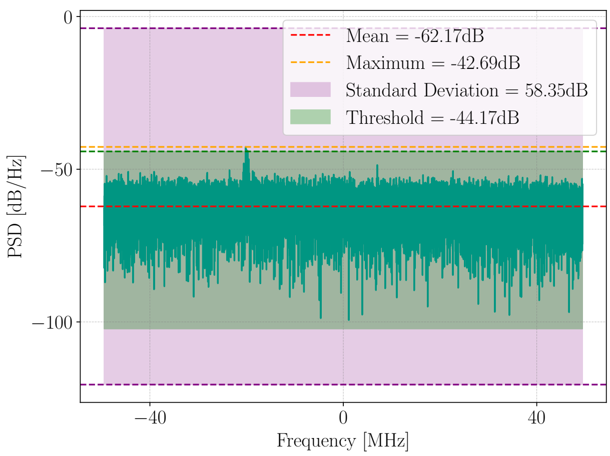

(c) Abdali, Kuwait

Figure 3.6: According PSD rangeline with maximum SIR for each quicklook image without visible RFIs but SIR profile with RFIs above threshold $T = 18\,\mathrm{dB}$.

3.2.2 Global Noise Scenes

Another data set used for the following work consists of globally distributed noise pulses. Before and after every imaged scene short 50-100 receive-only sequences measuring the system noise are recorded. This allows to annotate the noise in the processed image products. Those are not required to be mitigated since they are not part of the SAR focusing or image. However, they can be used to detect X-band RFIs all over the world in the data takes already acquired. As the noise pulses are short and available in a local data base this reduces the computational effort compared to the full level-0 data. Subsequently, there is no reason to depict images, have a look at all rangelines of the noise pulse or reduce the regarded bandwidth to the SAR signal's bandwidth to check the pure existence of an interferer at the according coordinate. Therefore, each noise pulse is regarded with the whole bandwidth in raw data and simplified to its maximum in the noise sequence per range sample only. So every noise sequence taken before a data take is summarized into the rangeline with the maximum power peak value only and computational effort is reduced.

3.3 Ground Truth Data Set

Ground truth is a term to describe the validated truth of the labelling of a data set. It is particularly important using DL methods. This is because there huge data sets are used in order to train a network for a task since its performance depends highly on the quality of the data set. Since the total amount of rangelines of the data set is in the order of 3 mio. RFI-free vs. 75 000 RFI-affected PSD profiles it is not feasible to check them all manually nor to quantify the RFIs adequately. Therefore, the SIR filtering is used in order to implement a ground truth basis with RFIs > $18\,\mathrm{dB}$ considered only. Furthermore, a subset of the full signal data set is taken. It consists of five files in total. Three of them are RFI-free and later on used for training. One of them is RFI-free and later on reserved for validation. The test set is a combination of RFI-free and RFI-contaminated files. It consists of two RFI scenes as depicted in Fig. 3.7 (b) and (c) in order to capture a diversity of RFIs and obtain representability and one RFI-free scene as depicted in Fig. 3.7 (a).

In order to ensure the ground truth in the following context, the SIR rangelines are taken as a detected baseline to compare the selected architectures performances to. Furthermore, only the additional detected RFI affected rangelines are checked by hand in order to evaluate the improvement from the benchmark of detecting strong RFIs with a SIR > $18\,\mathrm{dB}$. Therefore, a lower bound on the detectable RFI strength cannot be measured and neither does a full ground truth labelling of the scene exist. However, the ground truth of the training and validation set is validated manually. Also the RFI rangelines include clearly quantifiable RFIs above an SIR of $18\,\mathrm{dB}$. Therefore, the filter is describing a clear and quantifiable lower bound for considered RFIs in the test data set. This ensures a fair comparison of performances with a verified benchmark. The improvement in detection is evaluated relatively to this validated ground truth subset. It includes the SIR filtering as a first step to create the benchmark and a subsequent manual check for verification.

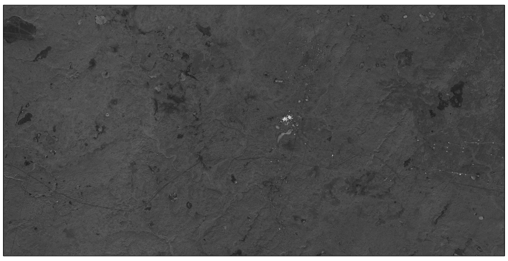

(a) RFI-free Alberta Canada

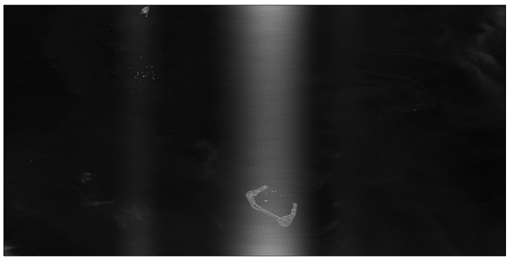

(b) RFI Thitu Reefs Pacific

(c) RFI Los Angeles USA

Figure 3.7: Quicklook scenes of the ground truth data set (a) interference free scene over Alberta, Canada; (b) and (c) scenes containing visible RFI over Reefs Pacific and Los Angeles, USA, respectively.

Tab. 3.1 shows an overview of the data sets introduced in this chapter. They are built and labelled as a basis for a fair and valid performance comparison. It has to be noted that the total RFI-scenes and samples are equal to the ones used for testing since they are used for testing only. Also the noise performance data set is a labelled subset of the total global noise data set in order to evaluate the detection performance on the noise data. Due to its large size the global noise data set cannot be labelled in a verified manner which is why it cannot be used for the performance comparison directly. Therefore, the noise performance data set as a labelled subset is introduced.

Table 3.1: Overview of all data sets used for this work including the complete data set, ground truth data set as a subset of it and the global noise data set for the performance test.

Attribute	**Complete data set**	**Ground truth data set**	**Noise performance data set**
RFI-contaminated scenes total and test	31	2	592
RFI-free scenes total	96	5	3984
RFI-free samples total	2 938 856	173 366	3984
RFI-free scenes test	93	1	798
RFI-free samples test	2 938 472	48 024	798
RFI-contaminated samples	74 754	4787	592
Total RFI-contaminated percentage	2.5%	2.7%	14.8%
Total RFI-free percentage	97.5%	97.3%	85.2%

Tab. 3.2 shows an overview of particularly the autoencoder's results on the aforementioned total global noise data set. The model is trained and validated on the same amount of samples as in the

performance comparison with the noise performance data set. In contrary to the noise performance data set however it consists of all TerraSAR-X noise sequences taken in stripmap mode from 2023 and 2024.

Table 3.2: Overview of the autoencoder's results on the extended total global noise data set. The model is trained and validated on the same amount of samples as in the performance comparison with the noise performance data set. In contrary to the noise performance data set however it consists of all TerraSAR-X noise sequences taken in stripmap mode from 2023 and 2024.

Attributes	**Global noise data set (results)**
Training scenes and samples	2390
Validation scenes and samples	798
Total scenes and samples	58 084
RFI-free scenes and samples detected	53 960
RFI-contaminated scenes and samples	1747
Total RFI-contaminated percentage	3.0%
Total RFI-free percentage	97.0%

4 Network Selection and Implementation

In order to analyze RFIs in X-band SAR data the RFIs first of all need to be detected. Since the RFI patterns in the images as well as in its PSD can vary a lot from, e.g., WBI, NBI to various WBI or various NBI, novel ML architectures are considered for RFI detection. Furthermore, since RFI-free PSDs on rangeline level are highly available in contrary to RFI samples themselves conventional anomaly detection methods are pre-selected as most promising and compared in the following work.

4.1 Data Pre-Processing

Feeding real, absolute SAR PSD profiles into a ML architectures obliges to some prior processing for correct training and detection. Therefore, in the raw data first of all the imaging data needs to be filtered from prologue and epilogue pulses transmitted before and after imaging pulses. Those include e.g. calibration pulses, synchronization pulses or receive-only (noise) pulses.

As depicted in Fig. 4.1 the resulting imaging data samples are fed into the pre-processng chain (a). Since they are still containing a Direct Current (DC) component at $0\,\mathrm{Hz}$ which could be misinterpreted as an RFI the data is first DC filtered, which is depicted in (b). Furthermore, the noise floor is cut off and the signal is truncated to the transmit bandwidth in order to only consider RFIs that are actually affecting the signal in the bandwidth relevant for the further processing steps. The result of this step is shown in (c). ML networks generally require the prior knowledge of two variables of the data set. One of them is the sample's shape. Therefore, the training, validation and complete testset is scanned for the maximum rangeline's shape meaning its samples and all other rangeline PSDs are interpolated to this maximum dimension in FD. This step is illustrated in (d), even though it cannot be seen in the plotted spectrum. However, the amount of samples and the oversampling factor simply increased for the maximum vector size handling of the PSD spectrum. The other variable that has to be known is the data sample's dynamic range so the ML network expects the data to be passed over in a normalized form. In order to fulfill this criterion (e) shows the normalization step mapping the rangelines PSD's values in a range of zero to one.

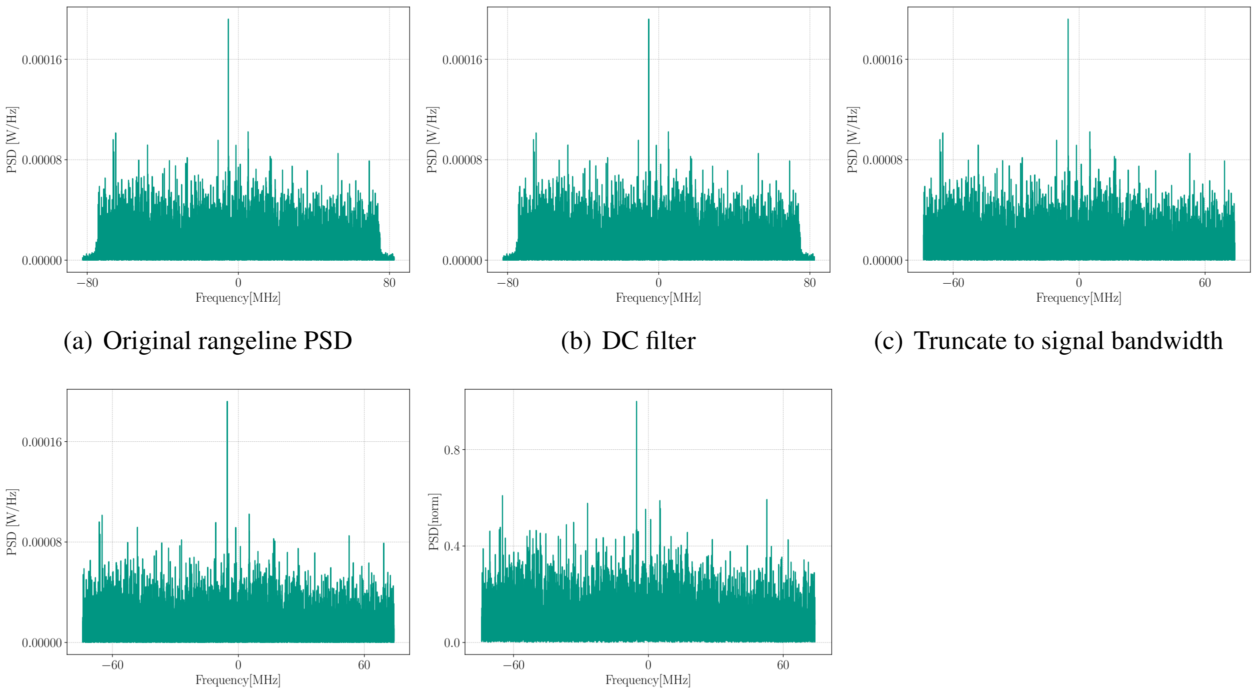

(a) Original rangeline PSD (b) DC filter (c) Truncate to signal bandwidth

(d) Interpolate (e) Normalize

Figure 4.1: Pre-processing steps applied to all input samples before inputting them into ML networks. The retrieved rangeline PSD in (a) is DC filtered in (b) and limited to the SAR signal's bandwidth in (c). In order to standardize for the ML network's proper function additionally each input sample is interpolated in (d) and normalized in (d).

4.2 Training and Validation Set Size

The standard approach for ML applications is to use $60\,\%$ of the data set for training, $20\,\%$ for validation and $20\,\%$ for testing. For anomaly detection only the normal data is used for training and validation and the separation here is implemented on full scene level. In order to achieve a general ground truth based on which the performance of the selected architectures can be compared fairly and to be universally valid and representative, a ground truth data set as a subset of the actual data set is used. It consists of five RFI-free scenes in total. Therefore, this leaves three RFI-free scenes for training, one for validation and one for testing.

Additionally, for testing the two RFI containing test scenes already shown in Chapter 3 are SIR filtered. The RFI rangelines are attached to the data set labelled as an RFI. In absolute data set sample sizes this means 123 871 rangeline samples are used for training, 1479 for validation and 52 811 for testing, consisting of 48 024 RFI-free test samples and 4787 containing RFIs. With an RFI ratio of $9\,\%$ to $91\,\%$ this is slightly increased to the complete data set that shows an RFI ratio of $2.5\,\%$ to $97.5\,\%$.

Fig. 4.2 exemplary shows the autoencoder's training history with the all samples of the ground truth training set consisting of three files in batches of 128 samples at a time such as its confusion matrix. Training in batches means splitting the full training set in a set of subsamples and training with this batched set before updating the weight matrix in the backpropagation step. This can be implemented manually or dynamically. Additional manual Batching in tensorflow is required for large data sets, when the full data set does not fit into the Random Access Memory (RAM)'s memory and dynamic batching is sufficient when the data set does fit in to the RAM's memory but not into the Graphics Processing Unit (GPU)'s. However, batch sizes typically vary in between 16 and 512. They are chosen based on the sample size and come with a trade-off. The lower the batch sizes the slower the training but the larger the batch sizes the more seldomly the weight matrix is updated which may affect the convergence of the model. It increases the risk to slow down optimization in each epoch and to run into a local, suboptimal minimum since the stochastical noise introduced by small batches is avoided. An epoch is the iteration of training with the full training set. With a small non-batched data set the whole training set is loaded into the model and the weights are updated after every epoch. However, with batched training the weights are actualized more regularly after every batch has been trained on in every epoch which maintains the favorable conditions for convergence.

However, reducing the complete training set in size to a training subset keeps the computational effort and duration when training as low as possible. Enabling a manual check of all used training and validation samples of the batch it even more importantly helps to ensure a ground truth. Therefore, the three training scenes are contributing to a training subset by 128 samples each, each sample referring to a rangeline or PSD profile. This results in a new reduced training set of 384 samples from three different scenes. Another 384 samples are taken from the validation scene yielding a validation set of also 384 samples.

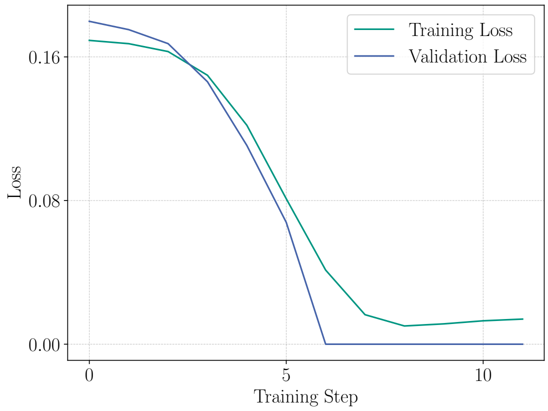

(a) Training history

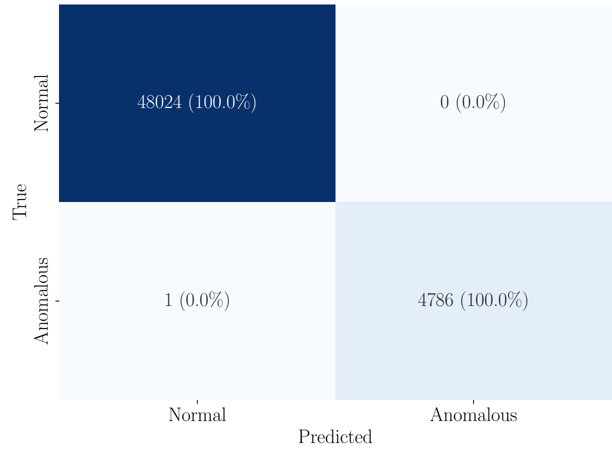

(b) Confusion matrix

Figure 4.2: The training history (a) and performance (b) of the autoencoder using all samples of the selected training and validation scenes of the ground truth data set is shown. The train data set is sliced in batches of 128 for training and it yields optimal performance meaning it is at least as performant as the underlying SIR filtering.

In comparison to training with the complete training set Fig. 4.3 exemplary shows the autoencoder's training history with the training subset such as its confusion matrix.

As can be seen, the performance is maintained on a constant level by increasing the number of training epochs. This means that the backpropagation, which is the updating of the model's weights according to loss optimization after each batch remains the same. However, the number of iterations over the whole data set is increased maintaining the same performance as feeding more data samples.

Since this enabled the ground truth for this subset to be validated manually and the performance is maintained, it is viewed as sufficient and is used for the following work.

Additionally, as can be seen from both confusion matrices, the resulting performance of the autoencoder yields a nearly optimal performance relatively to the SIR filtering. Therefore it is pre-set as a benchmark for the other pre-selected architectures.

4.3 Implementation and Performance Comparison

For the RFI detection three ML architectures are pre-selected and therefore compared regarding their performance. They were implemented in Python 3.10.8 using tensorflow for GPU and the keras library.

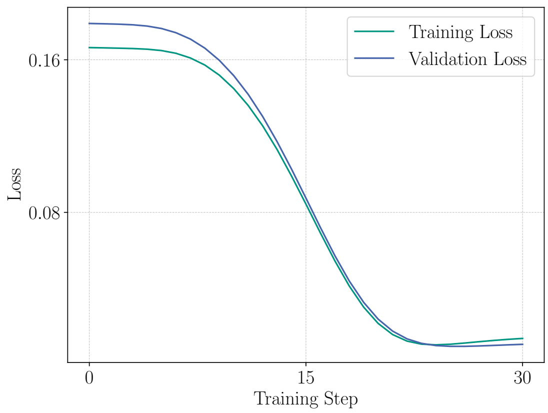

(a) Training history subset

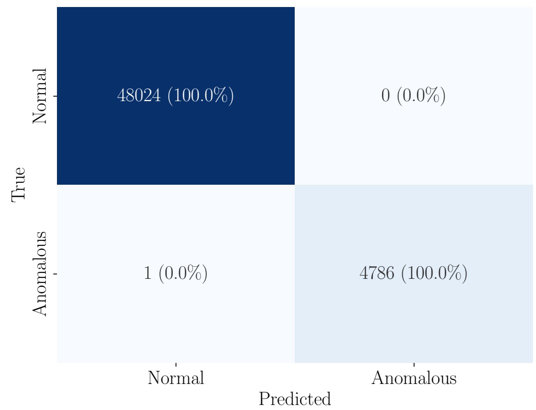

(b) Subset confusion matrix

Figure 4.3: The training history (a) and performance (b) of the autoencoder using only the subset of the three selected training scenes of the ground truth data set is shown. All samples of the one validation scene are used for validation. The train data set consists of 128 samples of each of the three training scenes that are manually checked and verified to be RFI-free. It still yields optimal performance meaning it is exactly as performant as training with the full training set size. Also it is still at least as performant as the underlying SIR filtering.

4.3.1 Isolation Forest

The iForest architecture is implemented as a package from the $sci-kitlearn\ ensemblelibrary$ and used as a pre-existing package. The features for the data set are of stochastical nature. Therefore the mean and maximum value which represent in the RFI case the signal and the RFIs PSD are fed as the principle components. However, in order to check for significance in other features additionally the median and the standard deviation σ are fed for training. The standard deviation's computation is shown in Eq. 4.1 with the number of datapoints n in the sample, the i-th data point of the input sample x_i and the arithmetic mean of the sample $\bar{x}$:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})}{n-1}}. \tag{4.1}$$

The median is computed as with the number of datapoints n in the sample and the n-th data point of the input sample x_n as

$$\text{Median} = \begin{cases} x_{\frac{n+1}{2}} & \text{if n = odd,} \\ \frac{1}{2}(x_{\frac{n}{2}} + x_{\frac{n}{2}+1}) & \text{if n = even.} \end{cases} \tag{4.2}$$

The *sci-kit learn iForest* as a model itself works in principle as described in the theoretical part in detail. It sets up an ensemble of decision trees splitting the data set according to random features at random values and stops splitting when a subset is isolated. Isolated subsets define outliers and therefore the sample is classified as an anomaly. The iForest can be fed by the attributes *contamination* which is by default set to $auto$. This assumes an anomaly ratio of $10\,\%$ from the original paper [LTZ09]. Another attribute used is *random state* which set to any integer ensures reproducibility by keeping the trained ensemble structure. Both of those are used for the simulation at hand. The unsupervised result as a scatterplot in Fig. 4.4 (a).

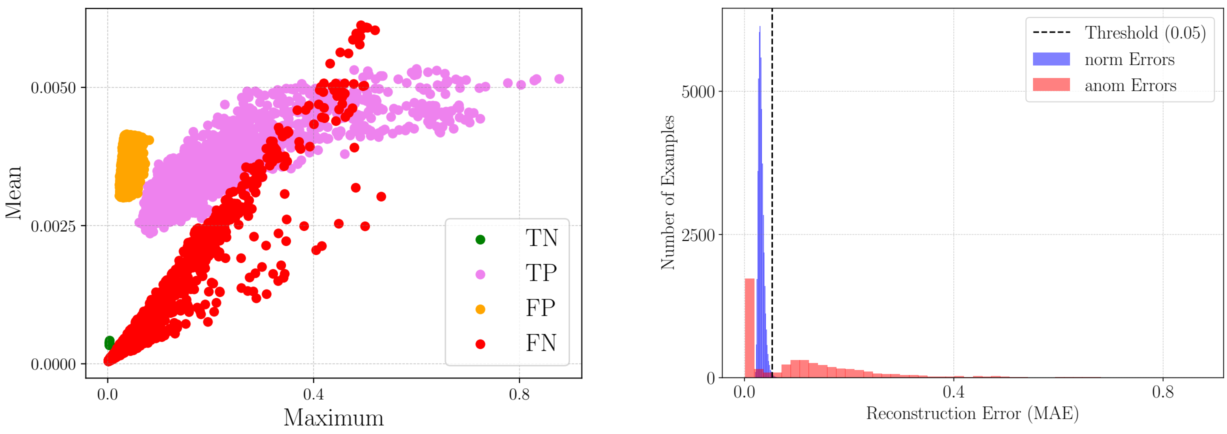

(a) Unsupervised scatterplot (b) Supervised error distribution

Figure 4.4: The iForest implementation is shown including the scatterplot over the mean and maximum of the samples as principle components after the unsupervised iForest stage itself (a). It shows that TNs and TPs cannot be separated uniquely from the TN and FP samples. This situation is similar for the MAE error distribution of all samples after the semi-supervised thresholding stage in (b) as the MAE for a part of MAEs is very low. This can be traced back to the mean and maximum value of each sample not being truely independent but rather correlated features in case of being RFI-contaminated.

As was verified empirically the maximum and mean are the principle components. The iForest itself is an unsupervised model. Having retrieved those principle components it can be transformed into a semi-supervised model by using the predictions for the principle different features and calculating the Maximum Absolute Error (MAE) as can be seen in Eq. 4.3.

The MAPE is computed by the difference of the element i-th predicted data point $\hat{y}$ in a sample to the i-th true data point $\hat{y}$. The maximum of all those differences can either be transformed into a percentage representation to the Maximum Absolute Percentage Error (MAPE) [Gil18]. However, it can also be taken directly as MAE per range line sample as it is implemented in the following work as

$$\text{MAE} = \max|\hat{y}_i - y_i|. \tag{4.3}$$

4.3.2 Autoencoder

The autoencoder as in detail described in Section 2.3 works by compressing the input data to a reduced dimension in the latent space. As a deep learning method it does so via various layers. The implemented architecture consists of an input layer, three hidden layers for compression in the encoder and symmetrically three hidden layers in the decoder before the output layer.

As dimensions for the latent space and the encoding and decoding layers a strong gradual compression of the input vector's size of 26958 down to 16, 8 and finally 4 data points in the latent space is chosen and verified empirically as the best performing option. This architecture's setup can be seen in Fig. 4.5.

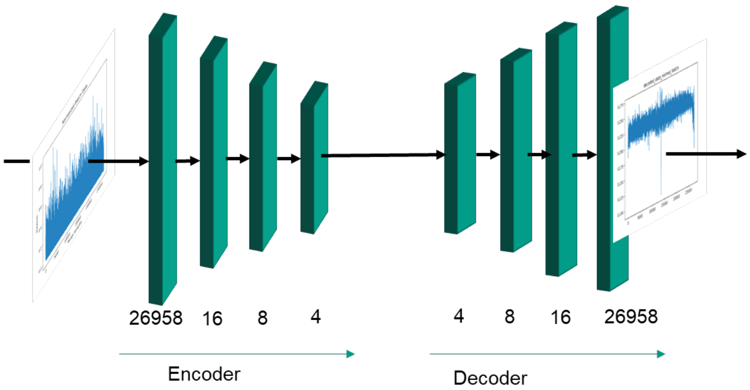

Figure 4.5: The autoencoder's network architecture consists of an encoder and a decoder each containing multiple layers. The implemented autoencoder has an input and output layer according to the largest input vector's shape of 26 958 data points per input sample that has to be known and standardized to all input samples. In the encoder this sample is compressed to vector shape's of 16, 8 and finally 4 in the latent dimension. All of these layers are chosen to be dense layers and the dimensions are determined to yield best results empirically. The decoder decompresses symmetrically. All layer's activation functions are chosen to be ReLU functions except for the last one that is chosen to be a sigmoid function since the decision anomaly or not is binary. During training the weights for each layer are optimized for the best reconstruction minimizing an MSE loss function. The autoencoder itself is a fully unsupervised model.

There is a variety of layers available as options to choose from. The most typical ones are dense layers, convolutional, pooling and normalization layers. While dense layer's perform the linear operation on the layer's input vector with the activation function as described in Section 2.3, convolutional layers work as filters selecting subsets of the whole layer's input vector at a time. They are also operating linearly by matrix multiplications. But they have another activation function attached at the layer's output which usually operates non-linearly. A pooling layer on the other hand effectively downsamples the data in between hidden layers passing on the prior valid information while reducing the number of operations required for the following layers. Normalization layers can be used at the input for feature scaling or as a hidden layer for batch normalization.

The hidden layers for the autoencoder at hand are chosen to be all dense layers and empirically verified to be sufficient for the RFI application. As activation functions for all layers the Rectifier Linear Unit (ReLU) activation function is used. They are the most popular in feedforward networks since they are computationally not too costly and accelerate the convergence of stochastic gradient descent. It basically computes the maximum of the input and zero [KK24] as can be seen:

$$f(z) = \max(0, z) = \begin{cases} z & \text{if } z \geq 0, \\ 0 & \text{if } z < 0. \end{cases} \tag{4.4}$$

For the last layer the sigmoid activation function is used. It is a binary activation function and therefor more applicable towards the binary decision at the output deciding if the sample is detected as an anomaly or not [KK24]:

$$f(z) = \frac{1}{1 + e^{-z}}. \tag{4.5}$$

The model is finally compiled using the Adaptive learning rate optimization algorithm (Adam) to optimize the learning rate for the $gradientdescent$ algorithm. In gradient descent there is a loss or error function that needs to be optimized. Here this loss is chosen to be the Mean Square Error (MSE). Compared to the MAE the MSE in this case performs best for training. It is computed by the difference of the i-th predicted data point $\hat{y}$ to the i-th true data point $\hat{y}$ squared. All of those differences are summed up and divided by the number of data points in order to get the mean and therefore the MSE per range line sample [Gil18]:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} (\hat{y}_i - y_i)^2. \tag{4.6}$$

The autoencoder itself is an unsupervised deep learning model. However, for anomaly detection its task is not only to learn how to reconstruct the input data. Instead, training only with normal RFI-free data, the autoencoder does optimize its abilities to reconstruct the normal data with an MSE reconstruction error as low as possible. Subsequently, testing on anomalous RFI data it is expected to yield reconstructions with a generally higher reconstruction error. This can be seen in Fig. 4.6.

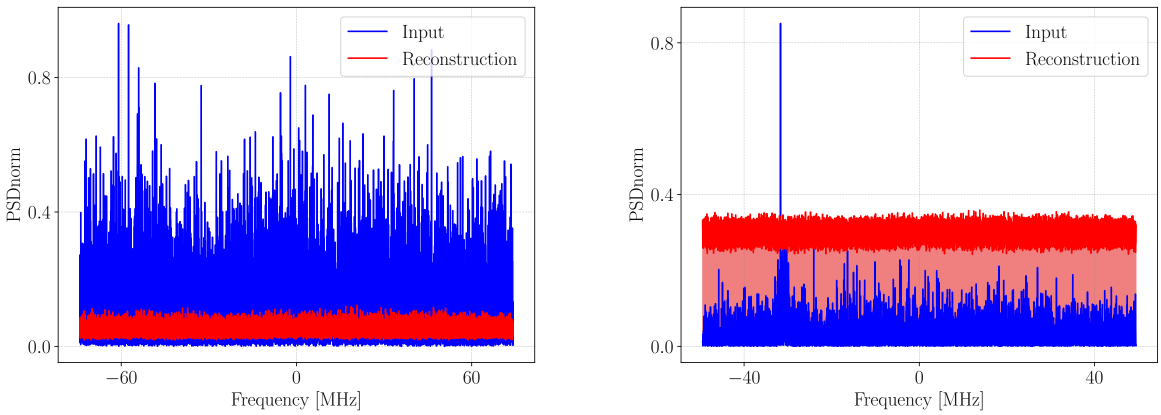

(a) Normal RFI-free (b) Anomalous RFI

Figure 4.6: An exemplary normal input sample (a) and anomalous input sample (b) of the autoencoder is depicted in direct comparison to the according reconstruction. The normal input sample's reconstruction lies within the area of the input samples themselves whereas the reconstruction of the anomalous sample is shifted on the y-axis. This is spanning an area of error within the input sample and its reconstruction.

These reconstruction errors cummulated in histograms are depicted in Fig. 2.9 (a). Separating them by a threshold transforms the algorithms functioning principle from an unsupervised learning to a semi-supervised working manner. The optimal threshold is chosen by sweeping it while checking the FPR and TNR with a labelled testset and the ROC curve as described in Section 2.3. This ROC curve is depicted in Fig. 4.7 (b).

Another measure for the performance that can be retrieved from the ROC curve is the Area Under Curve (AUC). It describes the area enclosed by the ROC curve as a first performance measure ranging from 0.0 to 1.0. i is an index for every point on the ROC curve that runs from 1 to the number of thresholds or points$n$ on the ROC curve as

$$AUC = \sum_{i=1}^{n-1} (FPR_{i+1} - FPR_i) \frac{TPR_{i+1} + TPR_i}{2}. \tag{4.7}$$

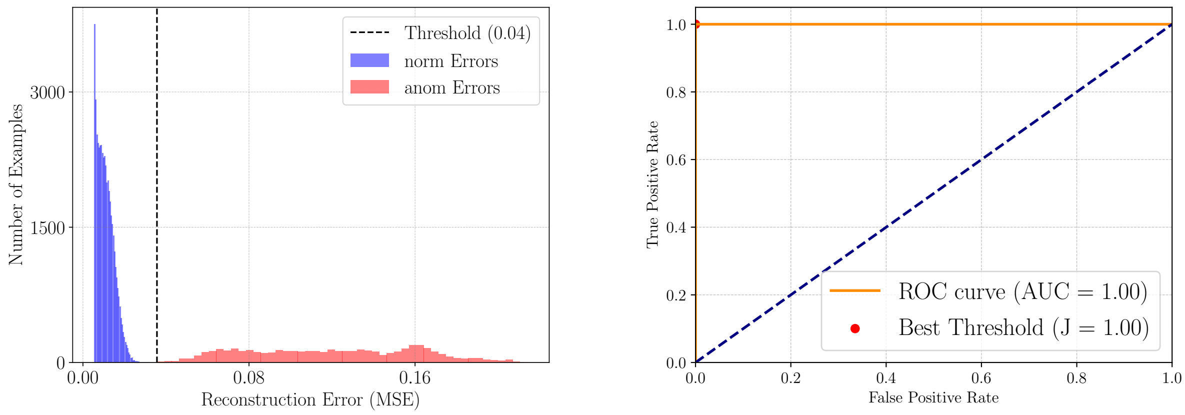

(a) Semi-supervised error distribution

(b) ROC curve

Figure 4.7: The autoencoder's reconstruction errors for the ground truth test set can be clearly seperated in their MSE distribution histograms (a). The optimal threshold is determined with the ROC curve that sweeps the threshold linearly and optimizes the ratio of the TPR and the FPR accordingly (b). This thresholding is based on the labelled test set and therefore supervised. This makes the combined anomaly detection of autoencoder and thresholding a semi-supervised technique.

4.3.3 Variational Autoencoder

The VAE's architecture setup is implemented analogue to the autoencoder. The main difference is the extension of the latent space to a distribution instead of a vector as can be seen in the depicted architecture in Fig. 2.10. This includes implementing the KL-loss described in detail in Section 2.3. The extension of the autoencoders architecture at hand to a VAE is based on the general architecture of the convolutional VAE in [Cho20].

However, just like for the autoencoder itself dense layers and the MSE as an error or loss function are used and the number of layers, such as their dimension remain the same. Accordingly, the subsequent semi-supervised steps with the thresholding techniques based on the ROC-curve remain the same. In Fig. 4.8 the distribution of reconstruction errors with the threshold optimized for the error distribution is also depicted for comparison reasons. As stated in Section 2.3 the threshold is swept and chosen based on the optimal TPR and FPR trade-off.

In order to compare the architectures regarding their performance, the evaluation metrics introduced in Section 2.3 and the confusion matrices in Fig. 4.9 are used. As can be seen in the confusion matrices of the iForest and VAE on the first glance none of their performances comes close to the nearly optimal performance seen in the autoencoders confusion matrix in Fig. 4.3 (b).

In detail Tab. 4.3.3 shows those evaluation metrics of the three architectures at hand compared.

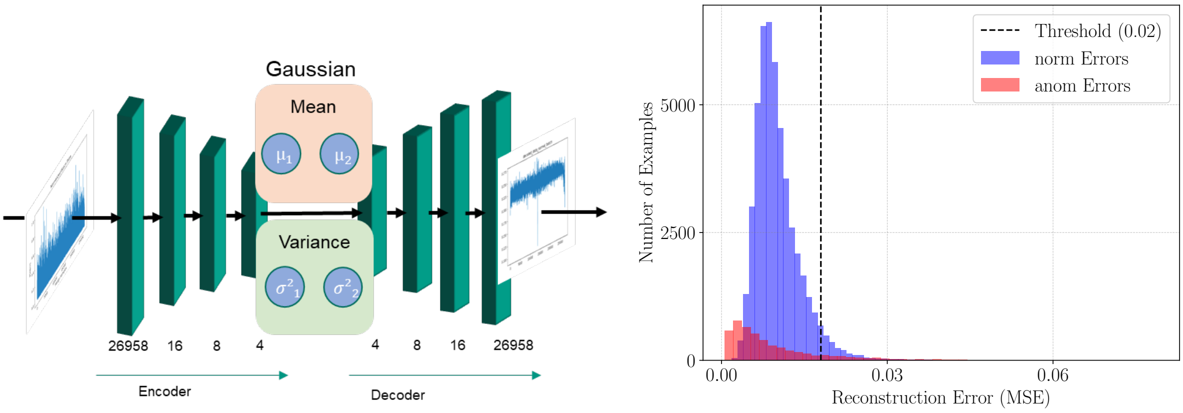

(a) Architecture (b) Supervised error distribution

Figure 4.8: The VAE's architecture is similar to the aforementioned autoencoder's. The three hidden compression layers to 16, 8 and 4 remain such as the dense layer type and the MSE loss function. They were emprically validated to still perform best. However, the input sample in latent space is mapped to a gaussian distribution here and represented as its mean and variance. The sample from this distribution randomly resampled as any point of this distribution (a). The following supervised stage calculates the MSE reconstruction errors for the normal and anomalous and determines the optimal threshold with the ROC curve according to the autoencoder (b). This also makes the VAE's anomaly detection semi-supervised.

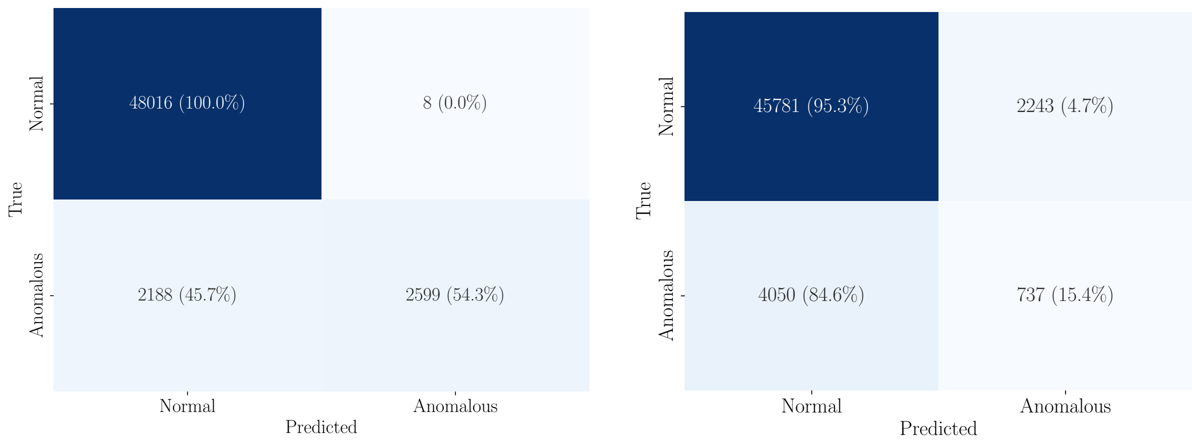

(a) iForest (b) VAE

Figure 4.9: The confusion matrices of the iForest (a) and the VAE (b) in direct comparison both showing performances significantly below the SIR filtering for RFIs with a SIR above 18 dB. Neither of them comes close to the performance of the autoencoder shown in Fig. 4.2 and Fig. 4.3 in Section 4.2.

Table 4.1: The comparison of the performance metrics for iForest, Autoencoder, and VAE is depicted including the AUC, best TPR, best FPR, accuracy, precision, recall, f1-score and the according best threshold.

Attribute	iForest	Autoencoder	VAE
AUC	0.61	1.00	0.38
Best TPR	0.59	1.00	0.17
Best FPR	0.00	0.00	0.06
Accuracy	0.9590	1.0000	0.8719
Precision	0.9344	1.0000	0.2219
Recall	0.5895	0.9998	0.1648
F1-score	0.7229	0.9999	0.1891
Threshold	0.052989	0.037976	0.015082

As can be seen from all the performance metrics the autoencoder performs significantly better than the other two considered architectures. The labelled data is set as the benchmark by SIR filtering with a threshold of $18\,\mathrm{dB}$. Therefore it has to be noted that everything that is below optimal performs worse than a classical SIR filter. Techniques underperforming this SIR filter for strong RFIs do not need to be further investigated since their sensitivity yields no improvement to traditional methods. While the iForest yields justifiable results regarding accuracy and precision, its recall is already below the desired behaviour. This means that the iForest's sensitivity and therefore its ability to identify all RFIs is poor. This is probably due to the iForest's simple architecture and particular suitability for data of high complexity meaning for data with various features or attributes. Since only the PSD per rangeline and some stochastic parameters retrieved from it are at hand to feed the iForest it is supposedly neither operating on the most suitable data set nor complex enough to detect the anomalies in the PSD profile per rangeline as the only important feature.

Even though being the most powerful and generalizing model the VAE shows the poorest performance. With an accuracy below $90\,\%$ and as well precision as recall below $25\,\%$ it is not only insufficient in detection. It still has a justifiable overall correction measured in the accuracy. However, its sensitivity measured in recall and accuracy of detected RFIs measured in the precision is underperforming significantly.

Having a look at its reconstructed samples in Fig. 4.10 leads to the conclusion that the VAE generalizes too much reconstructing the input sample. Therefore, the low reconstruction errors as well for normal but also for anomalous errors cannot be separated accordingly. A generalization so robust to fluctuations in the input would be very beneficial for the use case of creating new samples. Nevertheless, it makes the architecture unsuitable for the RFI detection case at hand.

It can be also noted that the autoencoders statistics on the ground truth data set being $100\,\%$ are nearly optimal or optimal taking into consideration minor stochastical effects. This means it is at least as performant as the underlying SIR filter of $18\,\mathrm{dB}$. However, it may also be impacted by a

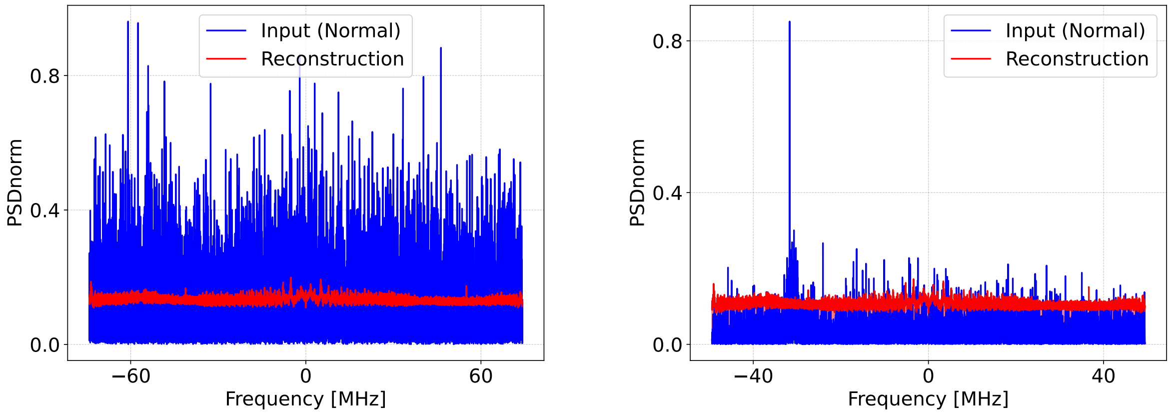

(a) Normal RFI-free (b) Anomalous RFI

Figure 4.10: An exemplary normal input sample (a) and anomalous input sample (b) of the VAE is depicted in direct comparison to the according reconstruction. The normal input sample's reconstruction lies within the area of the input samples themselves. In contrast to the autoencoder's reconstruction here so does the anomalous sample's reconstruction. This means there is no area of error spanning within the input sample and its reconstruction and the according MSE value is too low for a reliable separation.

limited difficulty of the subset. The ground truth data set is valuable as a fair and universally validated comparison basis of the three network architectures. However, including two RFI-affected scenes and one RFI-free scene only it is not solely applicable to evaluate a final realistic performance. Therefore another evaluation on the complete data set has to be implemented in order to represent more varying backgrounds and RFIs. This will be explained in the next chapter.

5 Autoencoder's Results and Extensions

In this chapter the selected architecture's performance evaluation is extended to the complete data set in order to achieve more realistic and representative results. Additionally, the autoencoder's ground truth performance is compared with conventional methods in order to make statements on the improvement of the novel detection method. Therefore, the SIR as well as the CA-CFAR filter is comparatively used. Furthermore, the signal data as well as the globally distributed noise data set is classified. It is analyzed regarding the occuring interferers and also a global interferer mapping in X-band w.r.t. the RFI's signal properties is presented. Finally an analysis of occuring errors is given.

5.1 Performance on the Full Testset

In order to make the evaluate a realistic performance a representative variety of RFIs and scene backgrounds needs to be considered. This goes beyond the ground truth data set used for the comparison of architectures. It needs to be tested on scenes including more varying backgrounds and RFIs. This is achieved by testing the model on the complete signal data set described in Chapter 3.

Fig. 5.1 shows the reconstruction errors of the autoencoder on the complete data set. As can be seen zooming into the first 100 number of examples in the histograms the separation is no more as perfect as in the ground truth example. This can be traced back to a higher variation and complexity in the RFI-free scenes and also in the types of RFIs in the RFI scenes. An optimal total threshold of 0.0276 for the MSE reconstruction errors is found which is set for the following analysis.

The full data set consists of 96 RFI-free scenes and 31 RFI scenes. The amount of rangelines in the complete test data set sums up to a total of 2 938 472 RFI-free rangeline samples and 74 754 RFI rangeline samples.

The RFI rangeline samples above the SIR of $18\,\mathrm{dB}$ are filtered for a labelled performance evaluation. This means also in the complete testset considering all RFI-affected scenes at hand only their rangelines that have an SIR of $18\,\mathrm{dB}$ or more are used as an RFI-affected test sample.

The training takes place with the same subset from Section 4.2 in order to ensure the manually validated ground truth for the model optimization. For the complete testset however a manual

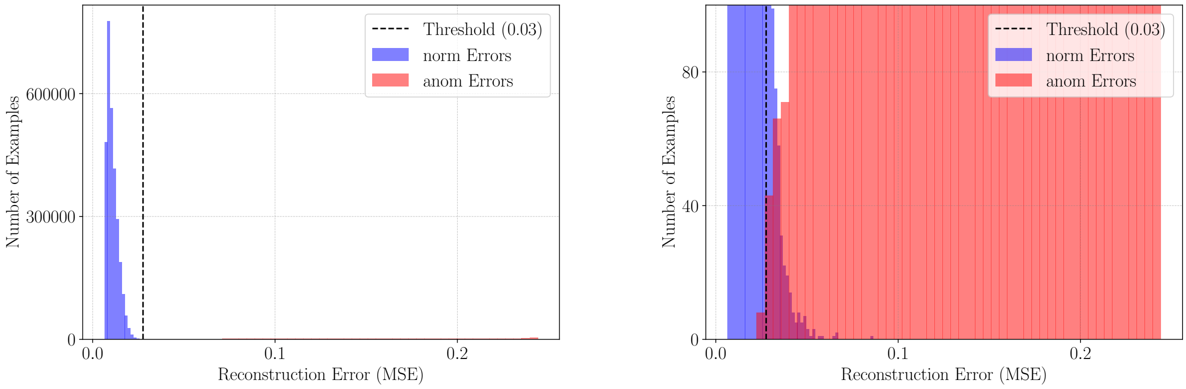

(a) Error distribution

(b) Zoom distribution with number of examples limited to 100

Figure 5.1: For a more realistic representation of variations in RFIs and scene backgrounds the autoencoder's performance evaluation is extended to the complete data set stated in Tab. 3.1. Its reconstruction errors and thresholds are depicted on a full-scale in (a) and also zoomed in to the first hundred samples in (b) as a detailed view of the overlap.

verification is not feasible. Even though the SIR threshold filtering works well for the RFI detection it sets a minimum below which potential weak RFIs remain undetected. These rangeline samples surrounding the SIR filtered RFI-affected rangeline samples in the RFI labelled scenes cannot be labelled clearly and would need to be checked manually. They are not included for the performance evaluation itself and therefore left out of the testset. This is because the majority of rangeline samples in an RFI-contaminated scene below the SIR threshold is RFI-free. However, the amount of samples below the threshold is in the order of several 10 000s. This makes their manual labelling extremely time consuming. Also it is not objectively quantifiable which makes it more representative to consider the rangelines with a SIR above $18\,\mathrm{dB}$ in the RFI affected scenes for the RFI-contaminated testset only. Subsequently, the remaining rangelines in the RFI affected scenes with a SIR below $18\,\mathrm{dB}$ are excluded from the data set and considered for the investigations at a later point in this work.

Actually a further step of the work is to even verify if RFIs below this or a lower SIR threshold have a visual impact on the resulting image or if they can be neglected because of disappearing in the received SAR echo. Also the RFI-free labelling is checked by the SIR filter since a manual verification and labelling of 3 mio. RFI-free rangeline samples is not feasible by hand but necessary in order to test the variety of monitored areas, backgrounds and targets.

Overall the full test data set with 2 938 472 RFI-free rangeline samples and 74 754 RFI rangeline samples after SIR filtering yields a ratio of $97.5\,\%$ to $2.5\,\%$. This strongly shows the applicability for anomaly detection.

As can be seen in the confusion matrix in Fig. 5.2 the imperfect separation in the reconstruction errors before has a very low impact due to the large data set and the probabilistic distribution. Therefore, the 746 FPs and the 12 FNs still yield a correct detection rate of $100.0\,\%$ rounding to the first decimal as it is implemented for the confusion matrix.

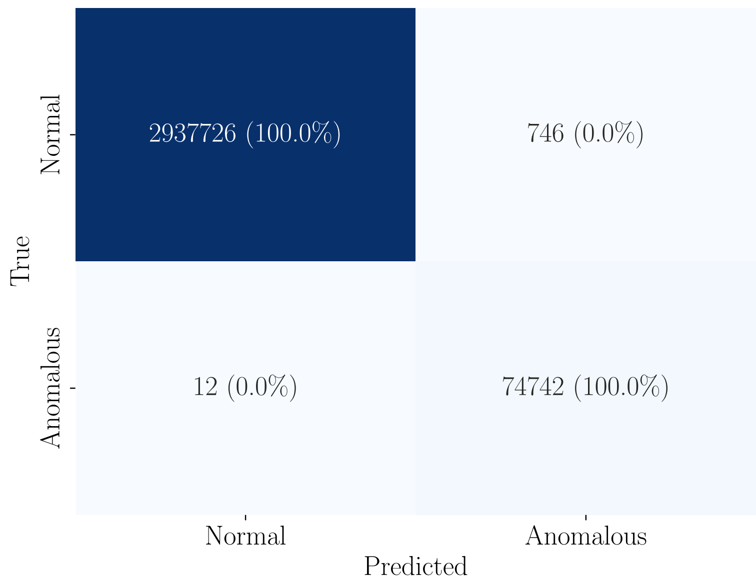

Figure 5.2: The confusion matrix of the autoencoder applied on the complete data set's test set consisting of all scenes is shown. It yields an optimal performance of $100.0\,\%$ relative to the underlying SIR filtering. This means the verified detection performance for RFIs with a SIR above $18\,\mathrm{dB}$ is a $100.0\,\%$.

5.2 Comparison with conventional Methods

The autoencoder as a novel deep learning method for the RFI detection in SAR raw data is selected due to its superior performance to other deep learning anomaly detection techniques as the iForest and the VAE. However, there are also conventional approaches targeting RFI detection and mitigation that have been established. Conventionally, they are used on simulated RFI affected data where the interferer meets theoretical conditions such as transmitting continuously and constantly in an isotropic manner. This differs from real RFIs since they do not meet those assumptions and increase the RFIs variance [FBKO23].

Therefore the autoencoder's detection performance needs to be held accountable against conventional methods in order to examine the achieved improvements. In order to do so, the conventional CA-CFAR which is commonly used in literature for RFI detection and mitigation across branches such as the automotive branch [Wan21] or passive remote sensing in radiometers [QCP$^+$21] is compared. Additionally, the autoencoder's detection rate is compared to the rather conservatively chosen SIR filtering with a threshold of $18\,\mathrm{dB}$ in order to measure a potential improvement. So far this threshold was used to determine the ground truth benchmark to evaluate the autoencoder's performance against. In the following it will still be used as a benchmark but the whole RFI-affected scene will be considered including the rangeline samples with a SIR below $18\,\mathrm{dB}$. That way potential improvements of the autoencoder or CA-CFAR relatively to the SIR filter can be evaluated.

The CA-CFAR is setting a threshold similar to performing a moving average over the rangeline from for the CUT from the left and right CFAR cells while excluding the directly neighboring guard cells padded by an application specific offset. This principle is explained in detail in Section 2.2.

As can be seen in Fig. 5.3 this principle of the CA-CFAR works as expected and well for NBIs with slim and high peaks. However, even strong WBIs remain undetected because the moving average is adapting to the continuous WBI pattern. This can be reduced by increasing the guard interval to very large numbers. However, this reduces the CA-CFAR behaviour of a moving average and adapting to the background noise to the simple SIR filtering approach that was used before.

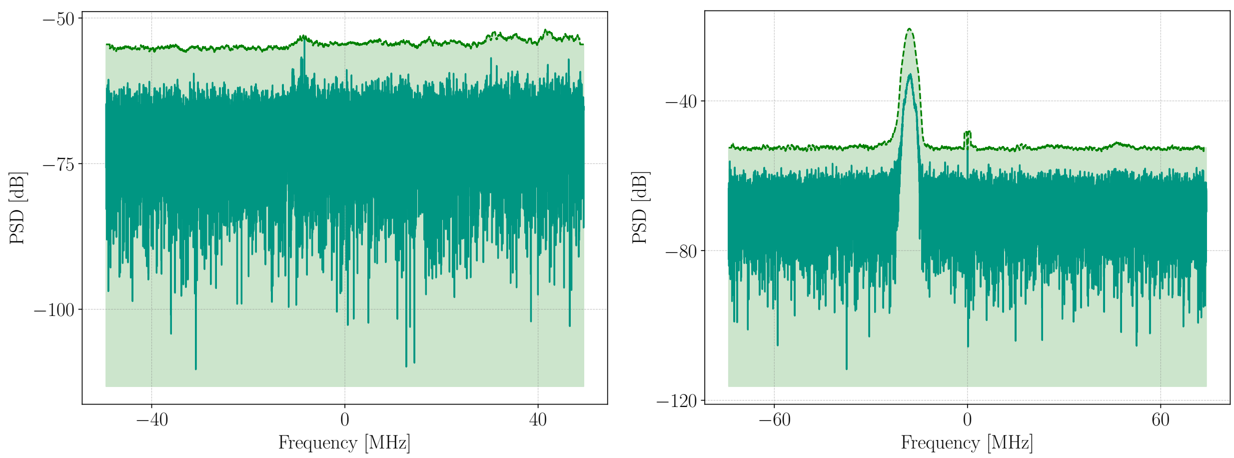

(a) NBI (b) WBI

Figure 5.3: The CA-CFAR threshold in exemplary rangeline PSD profiles are depicted. (a) shows the CFAR actually detecting the NBI by its peak whereas (b) shows that the WBI remains completely undetected and the detection is a case of luck because of a very narrow additional RFI in the spectrum besides the WBI. This can be traced back to the moving average that causes the CA-CFAR- threshold to adapt to the signal but also to graduate slopes as a WBI's artifacts.

In order to fairly compare this behaviour the two WBI scenes in the ground truth data set and one additional NBI scene are selected as a subset for testing. The exemplary rangeline samples in Fig. 5.3 are taken from those scenes. Tab. 5.1 shows the percentage of the total detected RFI-affected samples in the RFI-affected scenes. The total amount of RFI containing rangelines in the scenes cannot be determined since the manual load is too high. However, the SIR filtering with the threshold of $18\,\mathrm{dB}$ detecting all RFIs above that threshold marks the benchmark. The benchmark marks the traditional, conventionally achievable reference that the other methods are evaluated in relation to.

The percentage of the full scenes image that are RFIs with a SIR stronger than $18\,\mathrm{dB}$ is for the NBI scene $5.7\,\%$ and for the two WBI scenes together $8.7\,\%$. It is measured in rangeline samples. In comparison the CA-CFAR for the NBI scene detects an additional $0.6\,\%$ of manually verified RFI affected rangelines. However, for the two WBI scenes it detects only $0.67\,\%$ in total which yields a deterioration to the conservative SIR filtering of $-8.03\,\%$. Also it means that only $7.7\,\%$ of the strong RFIs with an RFI above $18\,\mathrm{dB}$ are detected.

Finally, the autoencoder for the NBI detects an additional $1.3\,\%$ of the image compared to the SIR benchmark which is also verified manually. Even more interestingly however, in case of the WBI scenes it even detects an additional $2.7\,\%$ in contrary to the deterioration of the CA-CFAR.

Table 5.1: The detected RFI percentages in the full WBI and NBI RFI scenes are depicted. A manual ground truth cannot be determined for all rangelines with a SIR below $18\,\mathrm{dB}$ due to the high data load. Therefore, the SIR filtering's percentage is taken as a reference or the benchmark and the conventional CA-CFAR as well as the suggested autoencoder are compared against it.

Detected RFI rangelines in scene	SIR	CA-CFAR	Autoencoder
NBI (one scene)	$5.70\,\%$	$6.30\,\%$	$7.00\,\%$
WBI (two scenes)	$8.70\,\%$	$0.67\,\%$	$11.40\,\%$

This result is highly promising for the autoencoder's improved detection rate of real world RFIs in real X-band SAR data that are not meeting the theoretical assumptions of simulators.

5.3 RFI Classification

Apart from directly mitigating the detected RFIs they can be also further analyzed and clustered according to certain characteristics. In the imaging data with 31 scenes a certain variation of RFIs may be represented but still only a limited variety of RFIs is pictured. Therefore, for the actual imaging scenes it is rather interesting to test mitigation techniques. However, using another data set namely the globally distributed noise data set described in Subsection 3.2.2 a way larger amount of RFIs can be detected all over the world. This is where proper classification or clustering according to selected features comes to play.

5.3.1 Imaging Data Set

As a first step it is rather interesting to locate the RFI in the PSD and to find the occupying bandwidth as forwarded information for further mitigation steps. This can be achieved by detecting the peak and mapping it to the RFIs center frequency f_c. For this sci-kit learns $find\_peaks$ function can be used to find local maxima. The bandwidth can be detected by its according power level. The straightforward approach would be to use the Half Power Beamwidth (HPBW) measure. However, this is found to leave significant RFI power artifacts remaining. Better results are achieved using the SIR threshold of $18\,\mathrm{dB}$ as a power threshold that is mapped to the according bandwidth. These results are even improved using the scene specific SIR threshold as an adaptive power threshold mapped to the according bandwidth.

This localization of one or potentially several separate RFIs along with their bandwidths is shown in Fig. 5.4. It is depicted along the three interestingly varying quicklook examples from Fig. 3.2 in Chapter 3. Connecting the dots from the quicklook images to their according PSDs some observations can be made.

Overall the brightness of the RFI pattern in the quicklook image can be translated to the power of the according RFI in the spectrum. The clear bright lines across the whole rangeline dimension as it is seen in the Los Angeles, USA (a) scene seem to connect to a strong but single and relatively narrow-band interferer as a peak. This RFI is assumed to result from the Los Angeles airports radar system. However they resemble the NBI from Section 2.1.3 that are common in land-mobile radar systems [TSHW19]. The hatched, wide sawtooth patterns observed in Liadong Bay, China seems to correspond to pulsed, peaky but distributed RFIs operating in several frequency bands within the SAR's imaging bandwidth (b). Since it is so close to a bay the guess of it resulting from maritime and harbour radar systems can be made. They look like the pulsed WBI from Subsection 2.1.3 that are common in ground-based radiolocation radars.

Finally, the sort of lines out of point target patterns including the semi-transparent hatched sawtooth pattern in Queensland, Australia is not so clear. It seems to result from a rather wide-band RFI at the edge of the SAR signal's bandwidth as it is seen in (c) but also in other rangelines from an NBI at a completely different frequency further within the signal's bandwidth. Generally speaking, WBI cause a rather wide but softer, semi-transparent stripe than the bright lines seen at the NBI. So potentially the rare pattern in the processed image is due to the the differing RFI artifact shapes and center frequencies during the data take itself as each rangeline is impacted at another point in time and at another location individually. Another factor could be the RFI having it's peak power just at the signal's bandwidth edge. Due to the matched filtering in the SAR processing this should decrease the RFI's effect in the focused image significantly. Therefore, these switches of RFI center frequencies from further within the signal's bandwidth and more narrow-band to further to the edge and more wide-band RFIs are likely to be the reason for the mixed RFI pattern in the processed mage.

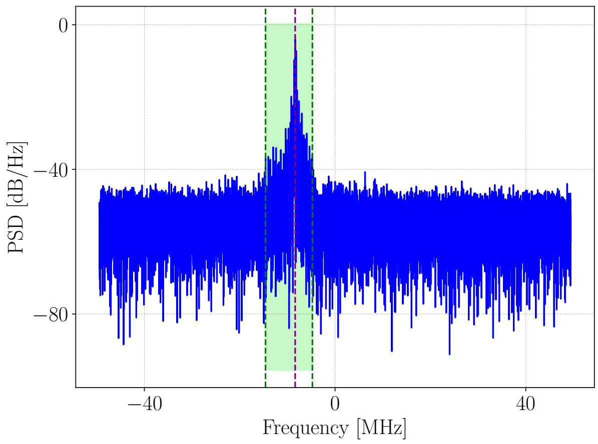

(a) Los Angeles, USA

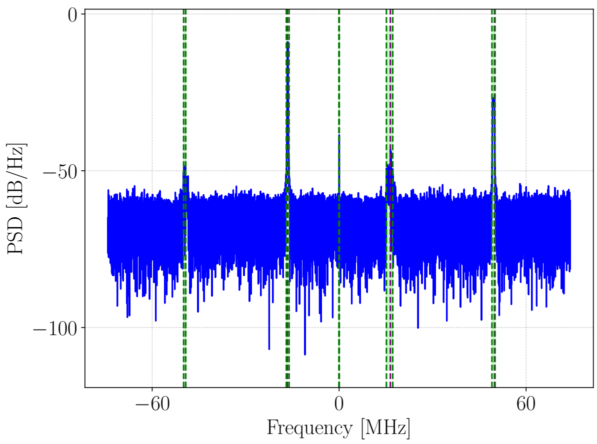

(b) Liadong Bay, China

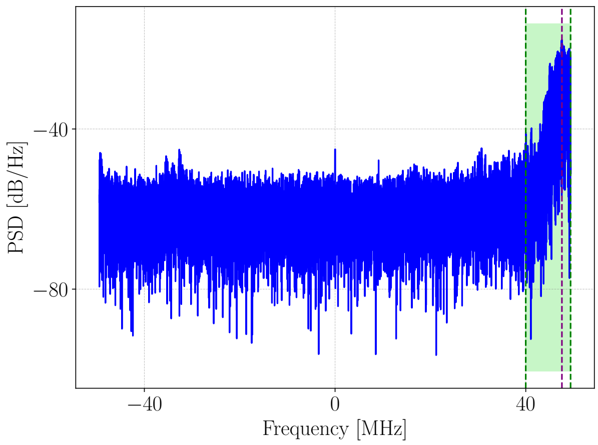

(c) Queensland, Australia

Figure 5.4: The corresponding RFI classifications to their quicklooks in Section 3.2 are depicted in their maximum's SIR rangeline PSD profiles.

5.3.2 Global Noise Data Set

Even though the location in the PSD with its bandwidth can be used as a first step for classification the actual use of the classification comes with the globally distributed noise data set. The global noise data set itself consists of all noise sequences taken in 2023 and 2024 by TerraSAR-X in stripmap mode. In order to evaluate the autoencoder's performance first however there is the noise performance data set created as a subset of this global noise data set as stated in Tab. 3.1 first. It is labelled by thresholding for training, setting a representative threshold and to evaluate the autoencoder's performance on the noise sequences. This is required since without the signal's characteristics the noise PSD profiles partially differ a non-neglectable amount. However, the test set used after confirming the performance with the subset contains the complete received noise pulses acquired by TerraSAR-X in the years 2023 and 2024. Here each noise scene is compressed into one rangeline containing the maximum in azimuth direction. This is done in order to keep the computational effort low and easily split into RFI and non-RFI scenes.

The training is implemented with $60\,\%$ of the RFI-free files of this subset. Those refer to 4788 training samples. There is no reduction to a subset feasible because noise pulses vary even stronger than the imaging scenes. This results in a testset of 798 RFI-free and 592 RFI rangeline samples yielding a ratio of $55.0\,\%$ vs. $43.0\,\%$ in the test set. It has to be noted that this results from the high data load required for training since the total ratio of the subset is $85.2\,\%$ to $14.8\,\%$. The higher percentage of RFIs here is a result of the limitation to 3 984 RFI-free scenes of the complete global noise data set for the performance evaluation.

As can be seen in the confusion matrix in Fig. 5.5 the noise performance data set yields a performance that is worse than in the signal scenes but for noise pulses still very acceptable. It measures a TNR of $99.9\,\%$ which means $99.9\,\%$ of truely RFI-free samples are detected as such. Also it shows a TPR of $95.7\,\%$ which means $95.7\,\%$ of truely RFI affected samples are detected as such. This yields an accuracy of $97.5\,\%$, a precision of $98.5\,\%$ and a recall of $95.7\,\%$.

For the classfication or clustering of the RFIs a number of features are selected. The first attribute is the location in the frequency band using the center frequency f_c. Its histogram is used informatively only instead of for classification. This can be seen in Fig. 5.6 (a). It is chosen not to be used for classification since the same interferer may be transmitting at various center frequencies at different times like the interferer in Queensland, Australia. This also opens a field for further investigation extenting the scope of this research.

Furthermore, the degree of occupied bandwidth is analyzed. In order to do so for each center frequency f_c the according bandwidth above the $T_{\mathrm{PSD,rl}}$ from Eq. 3.7 is determined. Several separated peaks with different f_c are set to have a minimum distance of 4000 samples to the next neighboring f_c, otherwise it is assumed to be the same RFI. Therefore, in case the power does not drop below $T_{\mathrm{PSD,rl}}$ within a one-sided window of 2000 samples the bandwidth is set to the according frequency f_c at the 2000-th sample. In the case of several separate interferers their bandwidths are summed

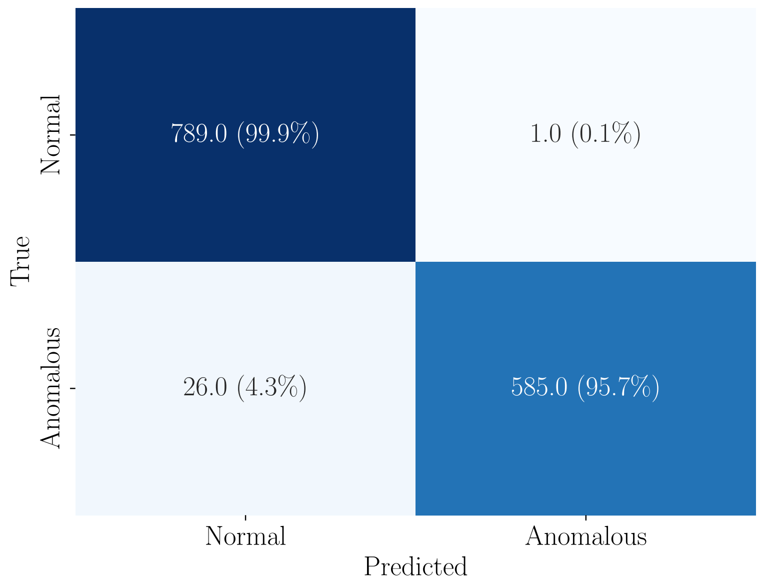

Figure 5.5: The confusion matrix of the noise performance data set from Tab. 3.1 is shown in order to evaluate the autoencoder's performance on a labelled subset of the global noise sequences. The full global noise data set is too large to be labelled as stated in Tab. 3.2.

up to a total occupied bandwidth that is used as a clustering metric. As can be seen in the histogram in Fig. 5.6 (b) the total occupied bandwidths are clustered into three classes. Narrow-band RFIs refer to $BW \leq 5\,\mathrm{MHz}$, wide-band RFIs to $5\,\mathrm{MHz} \leq BW \leq 20\,\mathrm{MHz}$ and very wide-band RFIs to $BW \geq 20\,\mathrm{MHz}$.

Another attribute chosen to characterize the RFI is its total RFI power in dB. The RFI's power is retrieved from the PSD spectrum by dividing the rangeline's PSD as a vector by the according rangeline's mean $\vec{\mu}_{\mathrm{PSD}_{rl}}$ which is the assumed signal in order to seperate the RFI's power from the combined signal's and RFI's power. Integrating over the RFI's bandwidth yields the RFI's power as

$$P_{\mathrm{RFI}_{rl}} = 10\log\left(\int_{f_1}^{f_2} \frac{\vec{\mathrm{PSD}}_{rl}}{\vec{\mu}_{\mathrm{PSD}_{rl}}}\,df\right). \tag{5.1}$$

It is used to analyze the RFIs overall strength transmitted over the whole spectrum. Another useful extension might be to consider the power normalized per occupied bandwidth. However, here the absolute total power is chosen in order to compare the RFIs strength directly. That way the impact of an increased power even if it is distributed over the whole spectrum and locally low per frequency can be considered and compared. As can be seen in Fig. 5.6 (c) the RFI's power is also clustered into three categories. A weak RFI corresponds to a total $P_{\mathrm{RFI}} \leq 20\,\mathrm{dB}$, strong RFIs range in $20\,\mathrm{dB} \leq P_{\mathrm{RFI}} \leq 32\,\mathrm{dB}$ and very strong RFIs are $P_{\mathrm{RFI}} \geq 32\,\mathrm{dB}$.

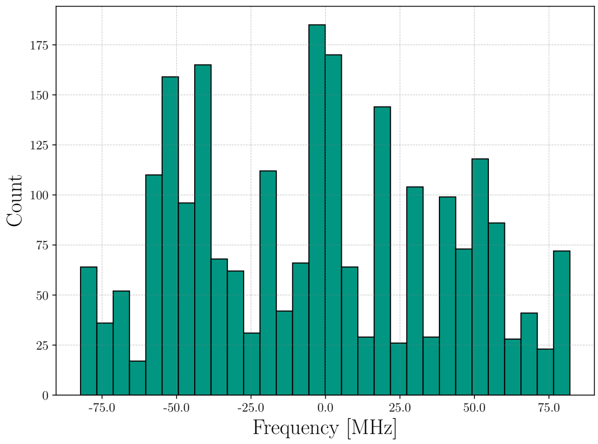

(a) Center frequencies

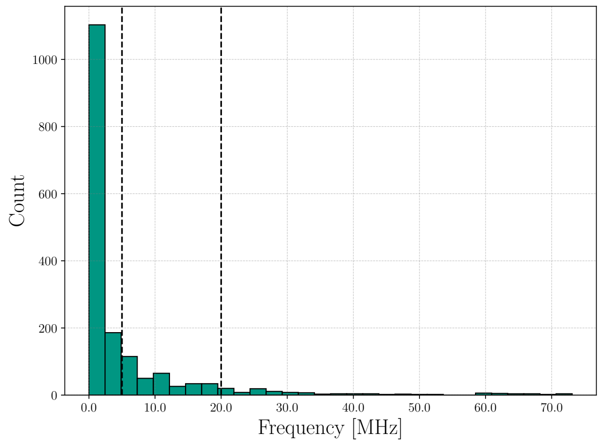

(b) Total occupied bandwidths

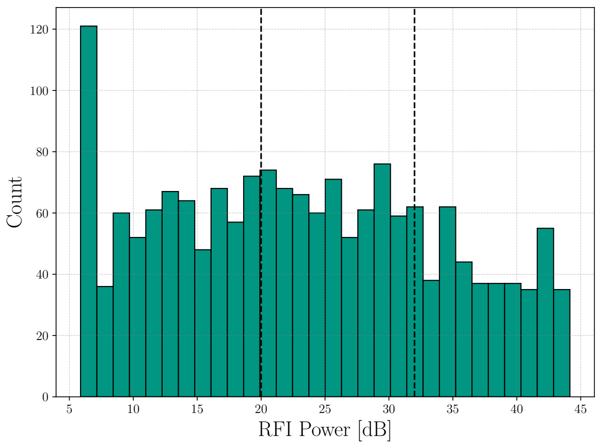

(c) RFI power

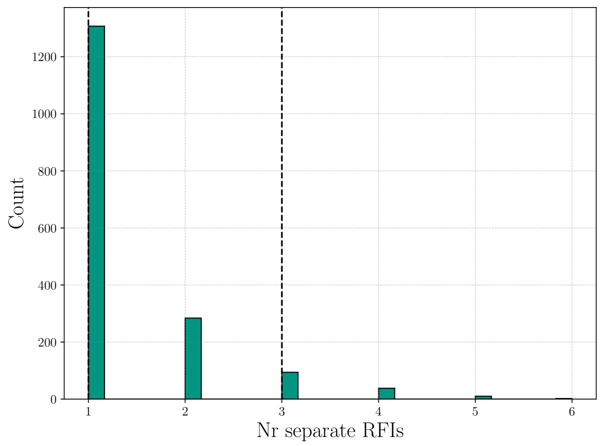

(d) Separate RFI count

Figure 5.6: The RFI's selected attributes and bounds for further classification of complete set of noise pulses from the years 2023 and 2024 are depicted. The classification is based on the features total occupied bandwidth, total power of the RFI(s) and count of separate RFI(s).

The last metric chosen corresponds closely to the separation of center frequencies f_c with the minimum distance of 4 000 samples set in between. By counting the peaks or center frequencies f_c determined per PSD profile or rangeline, the separate RFI count is determined. This corresponds to the chosen spectral distribution metric. As can be seen in Fig. 5.6 (d) the RFI count N_{RFI} is also separated in three classes. The single RFIs correspond to $N_{\mathrm{RFI}} = 1$ directly, $2 \leq N_{\mathrm{RFI}} \leq 3$ are summarized as distributed RFIs and $N_{\mathrm{RFI}} \geq 4$ are clustered as very distributed RFIs.

This classification is applied to the global noise data set in order to retrieve a global mapping and make statements on a global level. However, the data sets global distribution has to be taken into account before that in order to ensure its ability to be representative and estimate biases and its informative value.

In order to gain this estimation Fig. 5.7 shows the global distribution of all 58064 images and therefore considered noise pulses acquired by TerraSAR-X and TanDEM-X in stripmap mode in the years 2023 and 2024. The resolution of this map is $1°$ to $1°$ patches which correspond to the TanDEM-X DEM tile grid. Each patch is identified by the according data take's center coordinate.

The colour scale shows a ratio of RFI detected to RFI-free rangeline PSD profiles per data take, since an interferer does not necessarily transmit constantly and not even necessarily in the same frequency band every time the satellite passes by. Therefore, it may be that while in one overflight the area appears to be RFI-free at another time an RFI is detected. This ratio is depicted by the colorbar on the right varying from fully RFI-free in the blue case to only RFI affected detected in the red case with a continuous gradient. Additionally, the information of how many data takes at all have been acquired at this certain patch has to be made accessible. This is implemented by the amount of transparency at the particular patch. While no data takes correspond to zero, it increases to the maximum of data takes continuously as well.

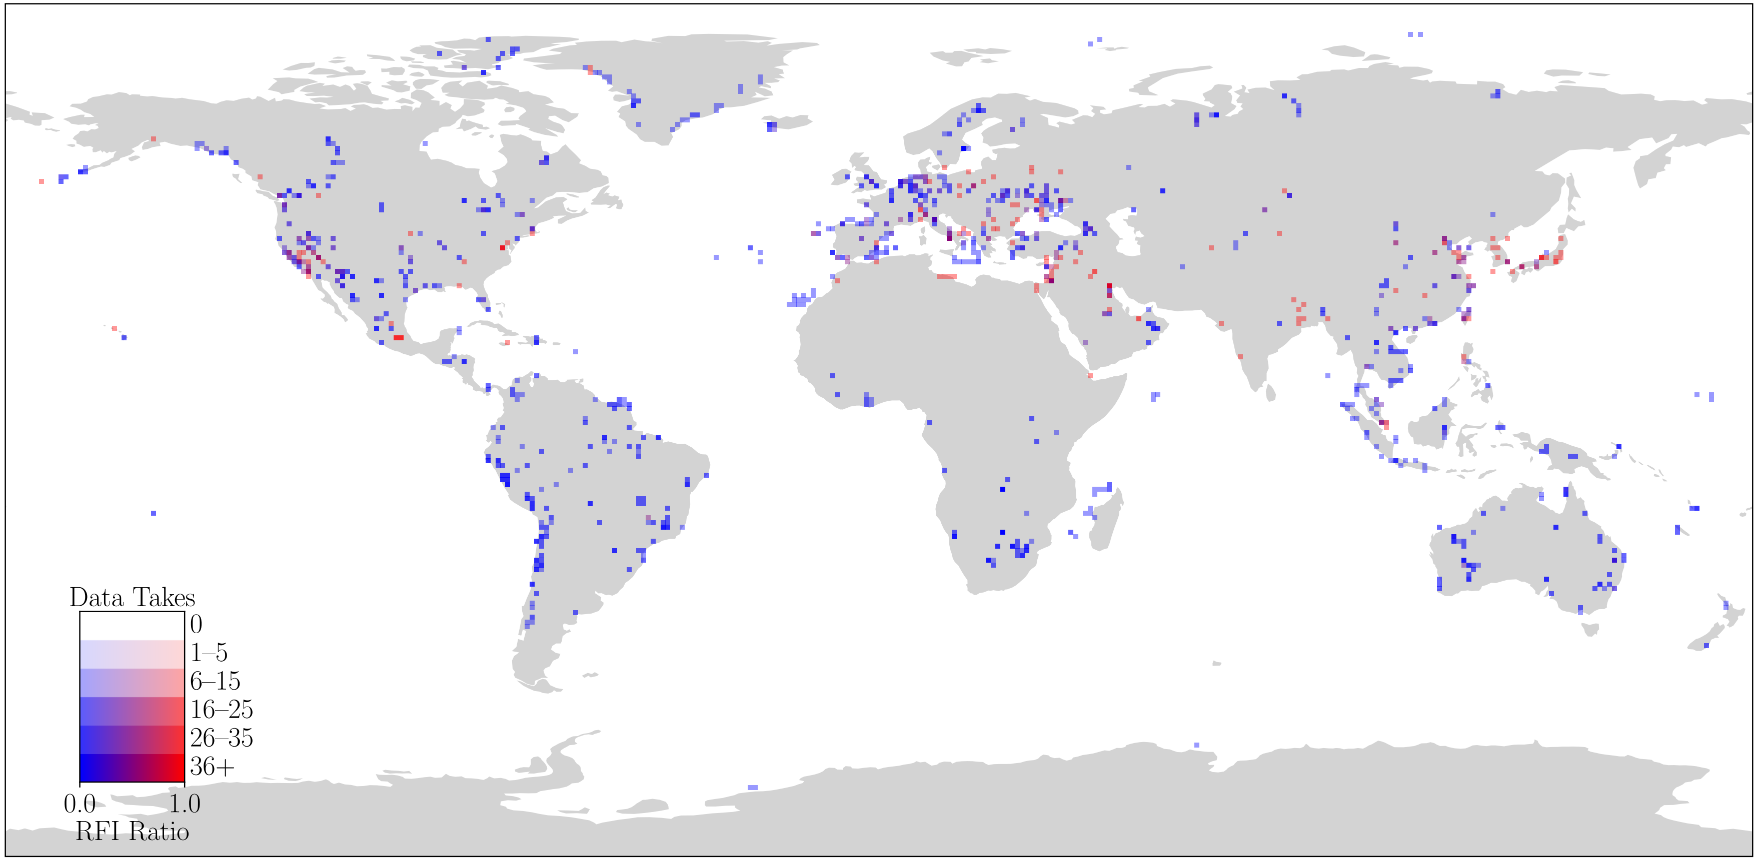

Figure 5.7: The global RFI ratio is depicted in a map visualizing all RFI-free and RFI-contaminated noise pulses of the TerraSAR-X mission's data takes taken in stripmap mode from 2023 and 2024. It visualizes the autoencoder's investigations in Tab. 3.2.

As can be seen from the global coverage of images and the RFI ratios, while having monitored the world on large-scale most of the RFIs seem to occur in prosperous and emerging market countries. Therefore, they especially cluster in Northern America, Europe, India and China. However, some more distributed RFIs within South America, Australia and Africa are also registered.

Having evaluated the global coverage of the noise pulses the globally detected RFIs can be classified in order to make statements about their geographical distribution. Firstly, Fig. 5.8 shows the global distribution of RFIs classified based on their total bandwidth occupation. While the strong occurence in Northern America, especially the USA, Europe and China is shown just as clearly as before it can also be said that most of the global RFIs can be classified as NBI. Most of the very wide-band RFIs can be found along the coast or within China, however also several wide-band RFIs can be registered in the US and Europe. The exact sources of those RFIs and reasons for this placement open up a field of further research.

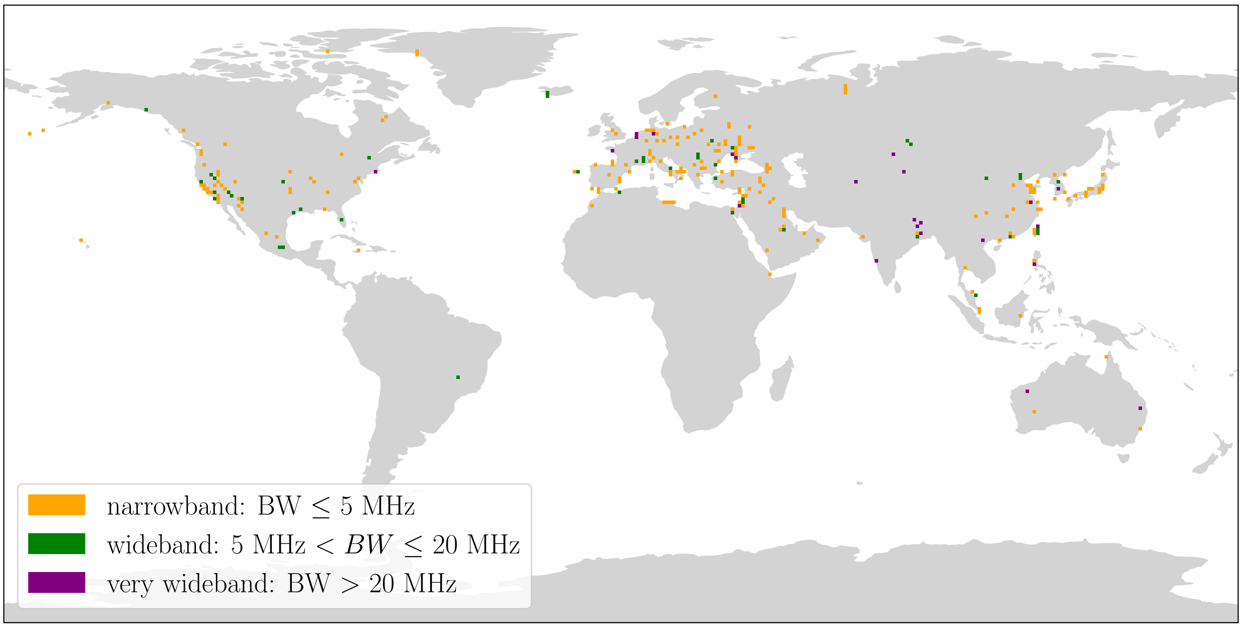

Figure 5.8: The globally detected RFIs of the TerraSAR-X mission's data takes taken in stripmap mode from 2023 and 2024 are visualized. Based on the autoencoder's results from Tab. 3.2 the detected RFIs are classified and depicted according to their total occupied bandwidth.

Nextly, having a look at the global RFI distribution according to their power classes, it can be noted that the very wide-band interferers also tend to be very strong ones especially within the asian continent. There can no clear dominance of strong or very strong interferers be noted. However, within the US mostly strong and very strong RFIs can be observed within the country and the weaker ones operate along the coast allowing the guess for the weaker ones to be maritime radars and stronger ones to be airport or military radar systems. However, this does not apply for Europe the Middle East and Asia, where strong and very strong RFIs can be just as well fount at the coast as weak RFIs can be found within the continent. This and the very strong interferers in the polar region as Greenland or Iceland might also enhance the guess that they might be rooted in air surveillance. However, for a more detailed analysis further research is required.

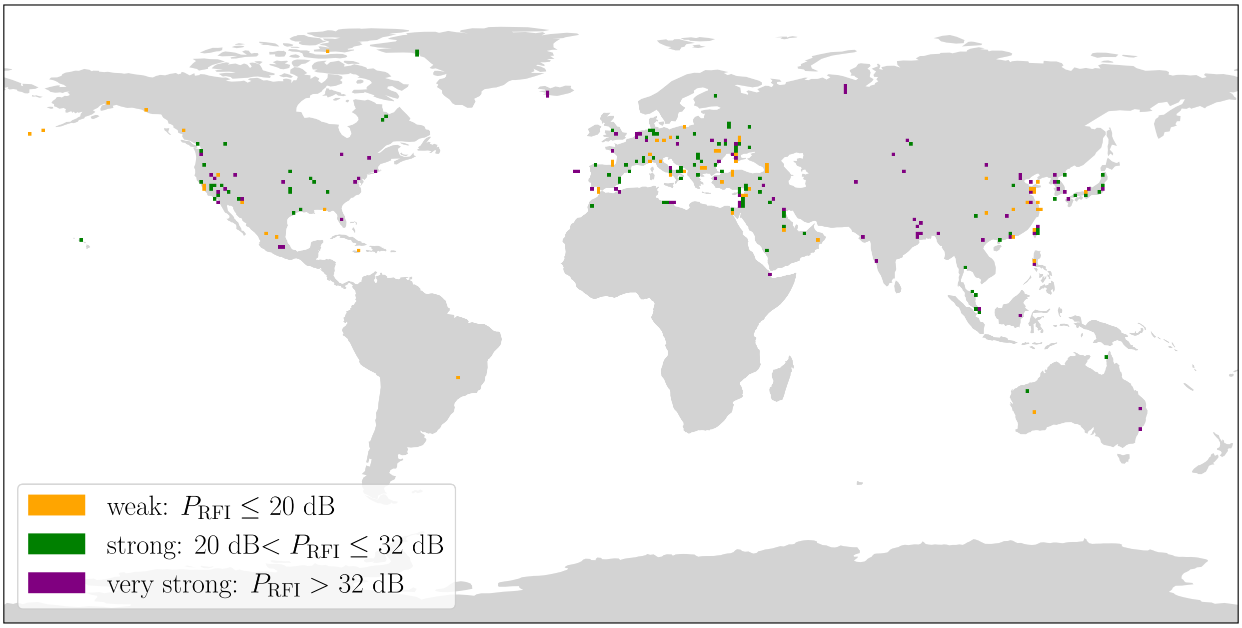

Figure 5.9: The globally detected RFIs of the TerraSAR-X mission's data takes taken in stripmap mode from 2023 and 2024 are visualized. Based on the autoencoder's results from Tab. 3.2 the detected RFIs are classified and depicted according to their total RFI's power.

Lastly, interpreting the spectral distribution of RFIs leads to a clear tendency again since the majority of detected RFIs appears to be single ones. However, plenty of distributed RFIs occur as well in Northern America, just as in Europe, preferedly at the coast and especially on the eastern Asian continent's pacific coast. Most of the very distributed RFIs correspond to the very strong RFIs that also tend to occupy very wide bandwidths within the Asian continent or along the Asian coast, particularly appearing to be on Chinese territory. This might be rooted in different regulation policies for airport or air surveillance systems of the according countries but in order to make more detailed statements further investigation is required.

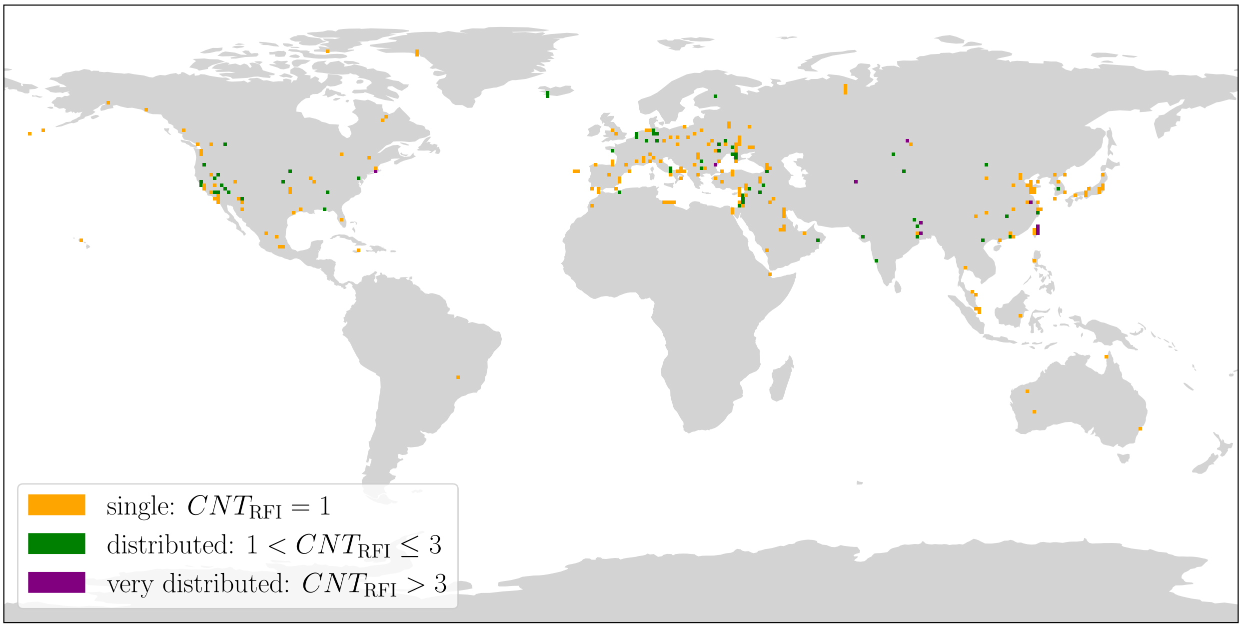

Figure 5.10: The globally detected RFIs of the TerraSAR-X mission's data takes taken in stripmap mode from 2023 and 2024 are visualized. Based on the autoencoder's results from Tab. 3.2 the detected RFIs are classified and depicted according to their total separate RFI(s) count.

6 Mitigation

So far a superior robustness of the autoencoder compared to the CA-CFAR and SIR filtering with respect to the RFI's bandwidth and shape such as detection performance in general are shown. However, the limits of the minimum strength of a detectable RFI is unknown. Mitigation is applied in order to remove the RFI from the imaging data while preserving the imaging information. The degree of mitigation depending on each detection rate is therefore evaluated in the following chapter.

6.1 Frequency Blanking on Rangeline Level

In order to do so, frequency blanking as the most basic form of mitigation is selected. The test mitigation of the selected scenes is implemented using an existing DLR internal software. The occupied bandwidth from Subsection 5.3.1 is characterized in each rangeline's PSD. For this their center frequency f_c and SIR threshold are chosen as the characterizing attributes. Frequency blanking sets this RFI bandwidth to zero so the containing information regarding power and phase is deleted.

The impact of this method differing NBIs and WBIs is evaluated on the NBI example scene and the two WBI example scenes introduced in Section 5.2. In addition to the visual comparison of the focused images for the evaluation of the RFI's impact reduction power profiles along the azimuth axis are used normalized to an RFI-free part of the image.

6.2 Narrow Band RFI

For the NBI the performance for the considered detection methods of SIR filtering, CA-CFAR and the autoencoder vary only little. Therefore, little to no visible improvements in the processed image are expected.

As Fig. 6.1 shows indeed the freuqency blanking mitigation yields very comparable results in the images. Mitigating on rangeline level the bright RFI lines in Fig. 6.1 (a) cannot only be mitigated in Fig. 6.1 (b), (c) and (d). This can be seen in the detail frame. But also underlying details like the ships are preserved. This shows that the underlying information is preserved for all three detection techniques.

(a) Raw data (b) SIR filter (c) CA-CFAR (d) Autoencoder

Figure 6.1: Comparison of the quicklook images of the mitigated NBI scene in comparison to the original unfiltered scene. After detecting with each detection method the data is mitigated by applying frequency blanking as a mitigation techique.

These conclusions retrieved visually from the processed images NBI scenes are reinforced looking at the power profiles along the azimuth axis. As can be seen in Fig. 6.2 (a) there is a considerable "hill" corresponding to the bright high power RFI lines in the original quicklook image with another two power peaks on top that correspond to the ships. As can be seen in Fig. 6.2 (b), (c) and (d) this basis hill introduced by the additional RFI power is reduced significantly mitigating all three method's detected RFIs while containing the ships' backscatter.

This mitigation performance yields nearly optimal results already for the benchmark SIR filtering. Therefore, the improved detection rate with an additional $0.7\,\%$ of the CA-CFAR as well as with an additional $1.3\,\%$ of the autoencoder do not yield considerable improvements on the mitigated and processed image's quality in the case of an NBI.

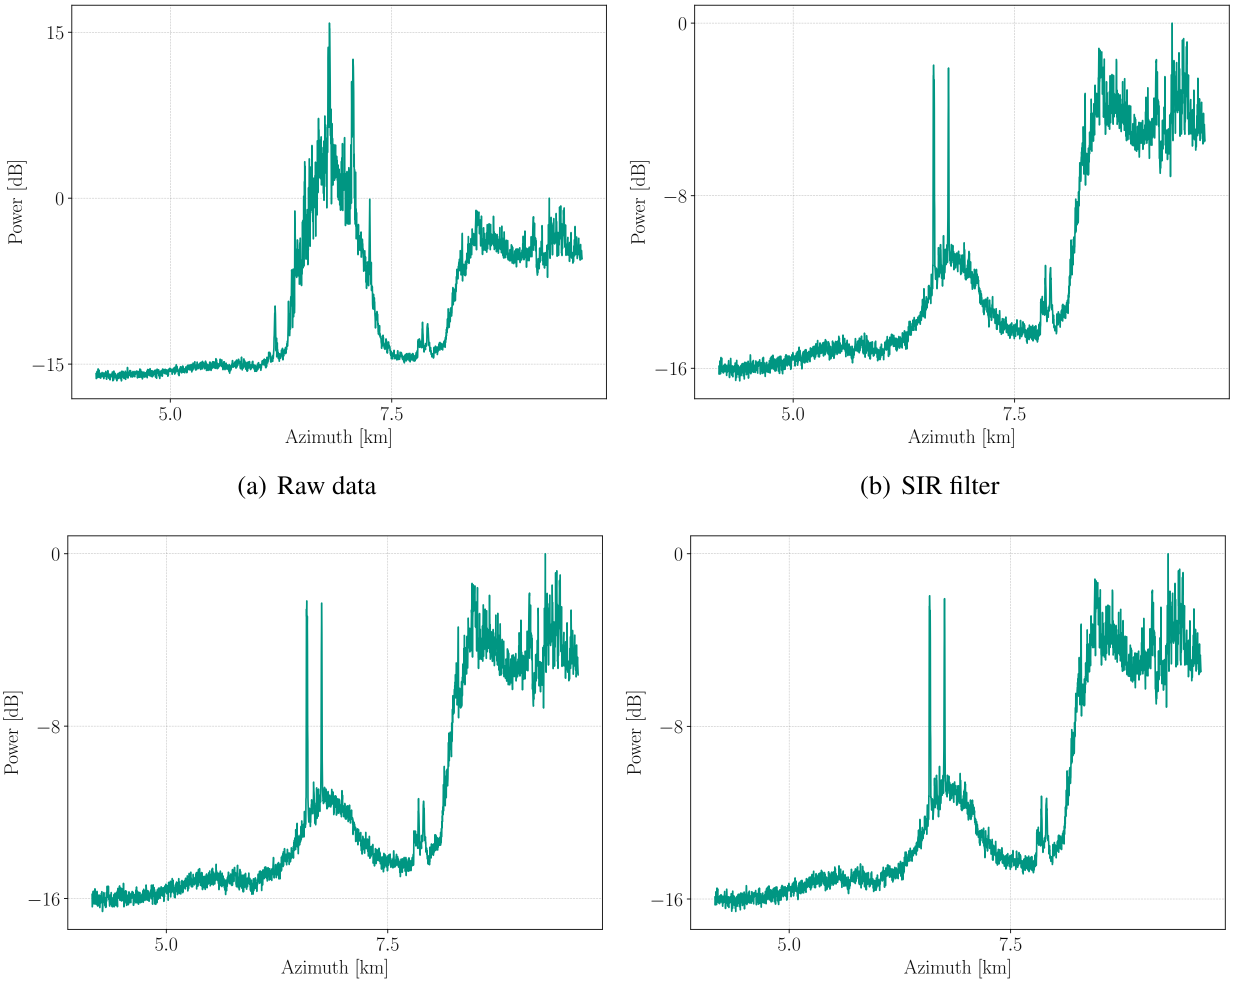

(a) Raw data (b) SIR filter (c) CA-CFAR (d) Autoencoder

Figure 6.2: Comparison of the power levels of the mitigated NBI scene in comparison to the original unfiltered scene. After detecting with each detection method the data is mitigated by applying frequency blanking as a mitigation techique.

6.3 Wide Band RFI

In the case of a WBI the CA-CFAR performs significantly inferior compared to the SIR filtering as the benchmark and also to the autoencoder regarding their detection performance. Verifying this detection performance's impact on the processed and mitigated images the CA-CFAR mitigated image is expected to contain significant RFI artifacts. Taking into account the detection performance and the NBI's mitigation results the SIR filtered image is expected to contain few impact of the RFI only with insignificant but slight improvements of the autoencoder regarding the mitigation performance itself. However, the scene specific SIR threshold has to be determined and verified manually. The same applies to the generalized SIR threshold of $18\,\mathrm{dB}$ for all considered RFI scenes in this work. It also has to at least verified to be suitable for the newly considered mitigation scene. The autoencoder however is trained on the rangelines' PSDs. Therefore, it is generalized and detects RFIs by their patterns in the according PSD profile. The resulting autoencoder's model trained on RFI-free data only and tested on the 31 RFI files can therefore be applied directly to new scenes without further adaptions.

Taking a look at the original WBI quicklook in Fig. 6.3 (a) it differs to the NBI pattern since the RFI pattern here is a wide and semitransparent stripe rather than narrow, bright lines following the SAR antenna pattern. As can be seen in Fig. 6.3 (c) the CA-CFARs detection performance of only $0.67\,\%$ of the whole scene shows an insignificant improvement to the original raw data image. However, the RFI stripe after mitigation using SIR filtering in Fig. 6.3 (b) and the autoencoder detection in Fig. 6.3 (d) yield comparable visual results. The image's part impacted by the RFI can still be made out, however it's impact is reduced almost completely containing some semitransparent narrow smeared lines along details like the ships along the rangeline axis only. Few but visually insignificant improvements from the autoencoder's mitigated image to the SIR filter can be noted.

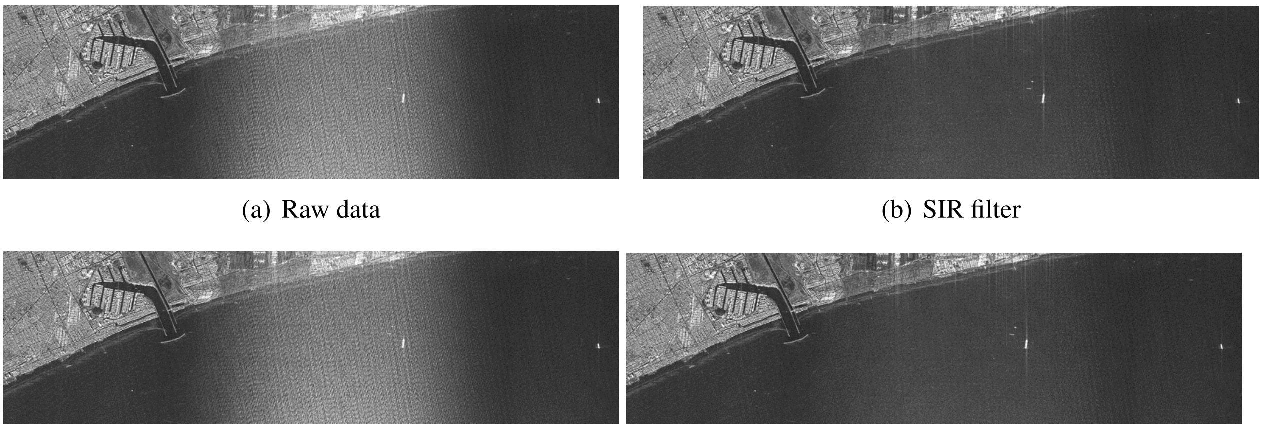

(a) Raw data

(b) SIR filter

(c) CA-CFAR

(d) Autoencoder

Figure 6.3: Comparison of the quicklook images of the mitigated WBI scene in comparison to the original unfiltered scene. After detecting with each detection method the data is mitigated by applying frequency blanking as a mitigation techique.

These observations and conclusions drawn visually from the quicklook evaluation is enforced looking at the power profiles along the azimuth axis in Fig. 6.4. The RFI power peak in the original image in Fig. 6.4 (a) remains for the CA-CFAR mitigated image in Fig. 6.4 (b). Accordingly, while the withdrawal of the RFI's power can be clearly seen as well in the SIR filtered mitigated profile in Fig. 6.4 (b) and the autoencoder's detected mitigated profile in Fig. 6.4 (d) the improved detection rate of $2.7\,\%$ of the latter cannot be noted in the power profile either.

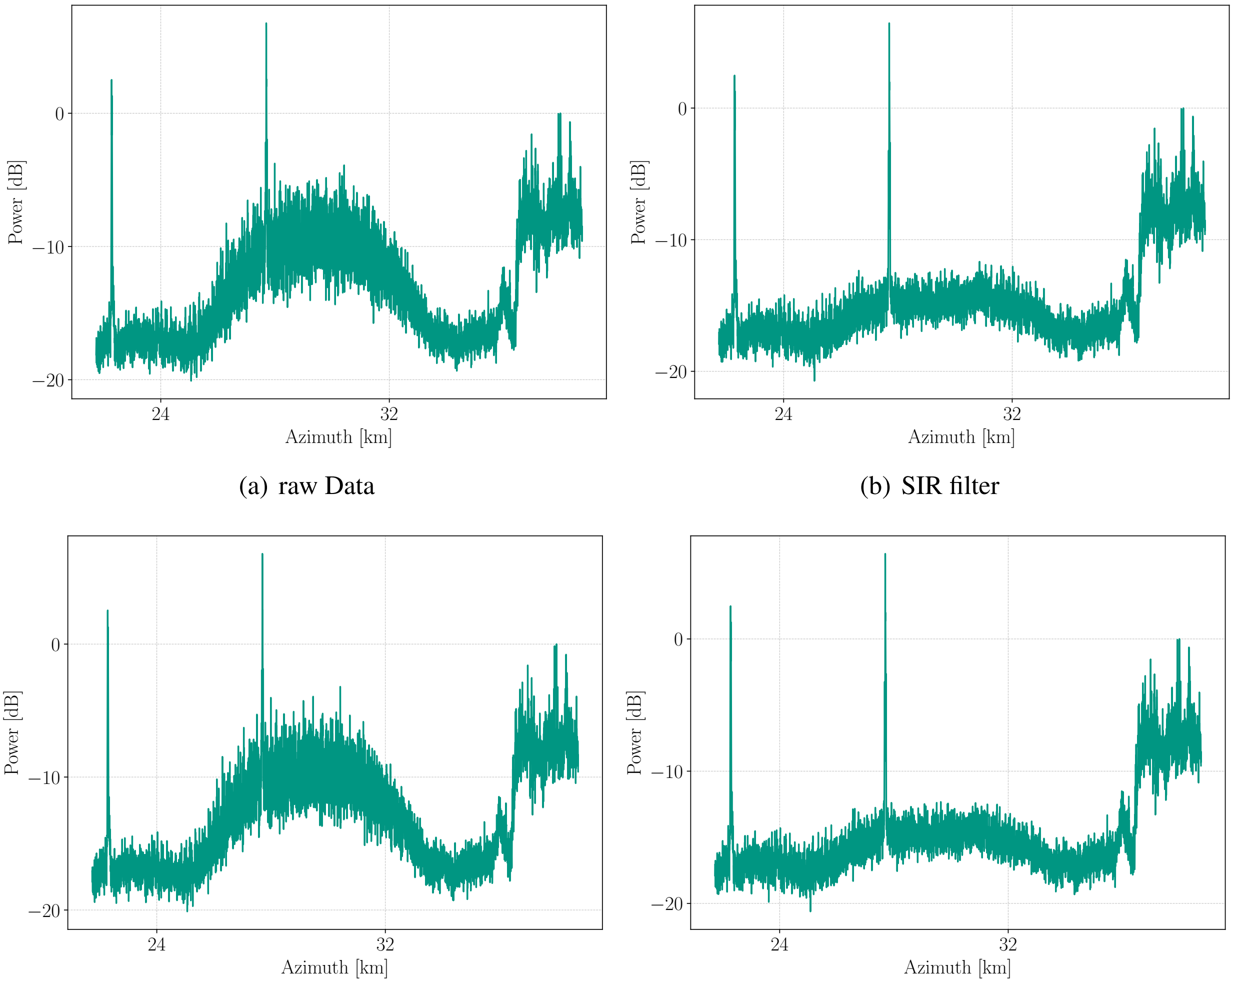

(a) raw Data

(b) SIR filter

(c) CA-CFAR

(d) Autoencoder

Figure 6.4: Comparison of the power levels of the mitigated WBI scene in comparison to the original unfiltered scene. After detecting with each detection method the data is mitigated by applying frequency blanking as a mitigation techique.

Therefore, while the autoencoder might introduce an enhanced detection rate and robustness to varying RFI patterns and backgrounds it does not seem to have considerably improved impact to the SIR filtering regarding the final processed images. However, in order to yield more holistic results the mitigation performance using alternative mitigation techniques needs to be investigated.

7 Discussion and Error Analysis

In order to complete the evaluation potential errors in the assumptions and their limits have to be considered. First of all weak RFIs below the SIR threshold of $18\,\mathrm{dB}$ might have remained undetected in the RFI-free data set. In the RFI containing scene data set those weak RFIs however are not included at all and therefore not considered for the optimization of the reconstruction error threshold and the following performance estimation.

Additionally, the conventional method of the CA-CFAR which was retrieved from literature especially used for the RFI detection and mitigation is also known to have a general weakness for wide-band targets. It does not seem to have impacted the RFI detection research a lot so far, which might be due to the focus on artificially simulated RFI affected SAR data which does not include these patterns within its assumptions. However, knowing that this weakness and data pattern exists an alternative version of the CFAR better suitable also for wide-band targets such as e.g. the OS-CFAR should be implemented and compared in further investigations.

It also has to be considered that the ground truth is only ensured for RFIs additionally detected to the SIR filtered quantifiable strong RFIs. Therefore, no absolute knowledge on the full RFI data set was obtained due to manual effort. Additionally, even with a complete manual labelling of a ground truth data set an objective and quantifiable lower RFI bound would have to be investigated. In the same stroke the minimum RFI strength that still has an impact on the image and still needs to be detectable as such would have to be set. Having found that even the additional weaker RFIs detected of the autoencoder do not significantly contribute to the optical quality of the processed image, it can be said that already the SIR threshold of $18\,\mathrm{dB}$ seems to be a valid approach for mitigation and the higher detection ratios use is more for classification. However, even for the ground truth data set this means it cannot be made an absolute statement on how many RFI affected rangelines all three methods might miss but the statement remains relative to the SIR filter as a benchmark.

Additionally, the ground truth scenes are selected randomly. While selecting two scenes decreases the chance to get only one particular RFI shape or degree of complexity of the scene it still does not ensure it to be representative. This can be seen in the comparison with the CFAR evaluating the performance on an NBI in an additional third scene. It also shows that both ground truth RFI scenes are showing rather WBI behavioral impact.

Generally thinking about a potential issue of optimizing the autoencoder to NBIs or WBIs it has to be noted that the autoencoder is trained on normal, i.e. RFI-free data only and therefore only adapting to RFI-free data. The shape of the RFI only comes into play when calculating the MSE reconstruction error. This one is used to optimize the threshold with the ROC-curve only. The optimal threshold of the complete test set including 31 globally distributed and varying interfering scenes is used. That is how a robustness towards varying RFI shapes is ensured. This can be seen in the widened but on the lowe-bound side limited error distribution of the reconstruction errors in Fig. 5.1 in section 5.1.

For classification as well as mitigation it has to be noted that the RFI's bandwidth is limited. It is detected until the SIR threshold or the maximum considered one-sided window size of 2000 samples as described in Subsection 5.3.2 only. Therefore, in some cases probably not the complete RFI power retrieved while in other there already may be retrieved some of the backscatter information. Meaning this general approach is not adaptive for each rangeline. For the imaging scenes it is adaptive for each scene, however for the noise scenes the general SIR threshold had to be used. Setting this threshold for peak detection as well means some faulty RFIs may be avoided, but some weak RFIs are also missed out on in classification which especially impacts the spectral distribution and RFI count classes.

Additionally, the absolute power is more likely to be high with several RFIs which seems logical. This comparison was intentional to be made, however the normalized power per Hz is an additional metric that should be combined with the absolute power in order to yield more robust statements in the comparison.

Overall it can be said that the classes are based on the test data set shown in the histograms. Therefore, they are limited by the variations and degree of being representative of as well the imaging RFI scenes data set such as the noise test data set's RFI scenes.

The use of this approach would be extending it to potential future standard direct detection and mitigation on a ground module for all downlink data. Regarding this extension it has to be noted that for the current model as well the input shape of the data has to be known such as the expected SIR threshold has to be computed for the scene. This means a general maximal limit to the rangeline dimension samples has to be set in order to interpolate smaller dimensions to this maximum. Another approach could be downsampling larger dimensions. Alternatively, using the SIR threshold a continuous computation and updating of the SIR threshold along the data take would require some further investigation.

8 Conclusion and Outlook

This work presents a holistic analysis of RFIs in general, such as their detection and mitigation using DL techniques in X-band SAR data. In order to do so a data set consisting of globally distributed scenes with varying backgrounds and types of RFIs in raw and imaging data was selected. It includes 28 RFI-contaminated scenes and 99 RFI-free scenes sorted based on a visual inspection of quicklooks on the imaging level.

A ground truth labelling method was introduced, derived from the PSD statistics of all selected RFI scenes and implementing a filtering of rangelines according to their simplified SIR. The SIR filter is detecting potential RFIs by the ratio between the signal as the rangeline's PSD mean value and the PSD's maximum value as the potential RFI's peak. From the logarithmic difference the absolute peak's SIR is calculated. With this simplification the RFIs' total power over the whole SAR system's imaging bandwidth is neglected. This is acceptable to filter strong RFIs but also leaves room for improveming to a finer detection especially for WBI. From that a statistically conservative SIR threshold of $18\,\mathrm{dB}$, relying on the standard deviation σ of the RFI scenes, was applied. This method enhanced the original labelling of clean scenes by incorporating weaker RFIs into the RFIs data set. The resulting data set on rangeline level includes 2 938 472 RFI-free (normal) rangeline PSD samples and 74 754 RFI-contaminated (anomalous) samples. This yields a total ratio of $97.52\,\%$ normal RFI-free samples to $2.48\,\%$ anomalous RFI-contaminated samples.

Furthermore, a global noise pulse data set was selected and constructed in order to create a global mapping of detected RFIs with the potentially enhanced detection rate. This can be used to contribute to the scientific community in the IEEE GRSS global RFI mapping for X-band that was depicted in Fig. 2.4 and retrieved from [GRS19].

In addition to those two data sets a ground truth data set as a subset of the complete imaging scenes data set was chosen. This was caused by the large amount of data in the whole data set that was too high for a manual verification of the ground truth of the labelling method. In the next step the potential improvements of the implemented methods were investigated. Therefore, a manually verified subset consisting of three RFI-free training files, one RFI-free validation file, one RFI-free and two RFI-contaminated test files was built.

For the anomaly detection, the ML algorithms of iForest, autoencoder, and VAE were pre-selected, implemented and optimized. This step contained preprocessing steps tailored for the real X-band SAR data. The algorithms were compared in terms of their performance on the ground truth data set. The retrieved results on the chosen evaluation metrics showed unambiguously that the iForest with an accuracy of $95.90\,\%$ achieves the largest reconstruction errors that lead to a relatively high threshold of 0.053. This suggests that it is a too simplistic approach for the data set and the input data's shape does not meet the iForest's most suitable use case's shape of having many different features.

The autoencoder with an accuracy of $100.00\,\%$ comes with smaller reconstruction errors and therefore also yields a smaller threshold of 0.038. This and the full set of evaluation metrics prove its suitability and a very promising performance, robustness and adaptability for the task at hand.

The similar VAE with an accuracy of $87.19\,\%$ shows the smallest reconstruction errors which results also in the smallest threshold of 0.015 overall. While this at first seems promising because it is very robust to fluctuations in the input data for anomaly detection it also means that it suffers from strong generalization. This leads to a suboptimal performance since the RFI-contaminated anomalies' reconstruction is too similar to the normal RFI-free one and a separation is harder. As a result, the autoencoder is selected for further observations as the most promising architecture for RFI detection in X-band SAR data.

When applied to the full imaging scenes test set, the autoencoder still achieves a $100.00\,\%$ accuracy, $100.00\,\%$ precision, $99.98\,\%$ recall, a best possible TPR of 1.00, and a best possible FPR of 0.00. This is even though there are misclassifications in the order of hundreds. However, this impact is neglectable relatively to the large complete data set size. While this is confirming the autoencoder's near-optimal performance it also suggests that it might be also an overly powerful and therefore overachieving tool for the RFI detection on rangeline level.

Therefore, a comparison with conventional RFI detection methods, such as CA-CFAR and the SIR labelling is drawn to evaluate potential improvements to the current state of the art. This comparison showed that the autoencoder outperforms both other methods for NBI as well as WBI. In the case of the NBI, the autoencoder improved by an additional $0.70\,\%$ to the second best performing CA-CFAR, while for the WBI scenes it increased the number of rangeline samples detected by even $2.70\,\%$ to the second best performing SIR filter. The CA-CFAR is found to be particularly weak in detecting WBI, which is however also a known weakness of the CA-CFAR.

In order to evaluate the benefits and necessity of the enhanced detection rate of the autoencoder the mitigation of the compared test subset for the comparison is demonstrated. As a mitigation technique frequency blanking is introduced as the simplest but exemplary method. While NBI mitigation techniques show comparable results across methods, WBI mitigation also confirms the extremely poor performance of CA-CFAR. However, also the autoencoder's significant impact on the detection rate's improvement turns out to be marginal in mitigation. Therefore for the actual mitigation of the image, the SIR filter derived from the data proves to be sufficient. However, it could be investigated if optimizing the threshold for the RFIs detected bandwidth that is to be blanked may yield more improved effects than an enhanced detection rate or can be combined with the latter.

What the enhanced detection rate does enable and does represent a benefit for is for global RFI tracking and mapping. This is yielded from the global noise data set retrieved from the noise pulses of all the globally taken TerraSAR-X imaging data takes of 2023 and 2024. Therefore, the global coverage is depicted as the ratio between RFI-free data takes and RFI-contaminated data takes on an 1° x 1°- resolution map taking into account the total amount of images taken in this area. Furthermore, those RFIs are classified based on their carrier frequency f_{c}, total bandwidth BW, power P_{RFI}, and individual RFI count N_{RFI}. Preliminary analysis of raw data RFIs links them to potential sources, with strong and wide-band RFIs concentrated in regions such as China and along coastal areas. Bright, slim lines are connected to single NBI, semi-transparent wide stripes to WBI, and sawtooth-like patterns in quicklook images are associated to pulsed, spectrally distributed interferers. Other patterns may be connected to interferers that change their center frequency f_{c} and/or bandwidth during SAR flight observations.

Further research could investigate the lower bound of visible RFIs, for example with the potential to enhance the SIR filter by incorporating RFI power or energy.

Additionally, DL based classification or clustering methods for the RFIs in the global noise data set could be implemented with simulated RFI types used for training.

Another field of further research could be exploring the OS-CFAR as an alternative conventional method to the CA-CFAR for RFI detection in order to cover the CFAR as a method adaptable for WBI as well as NBI and reduce known weaknesses and biases.

On a larger scale further investigations on the shape of RFIs in raw and imaging data according to their actual sources is required. Therefore, the raw data RFI shapes could be mapped more clearly with particular examples to typical patterns in the processed images such as their occurence in maritime and harbor systems, airport surveillance, air surveillance, military systems or other SAR systems. This extension of the classification and analysis offers valuable opportunities.

Finally, implementing a time-series analysis could provide insights into interferers switching center frequencies, bandwidths and maybe even their full operation mode over time.

Bibliography

[Air25] Airbus, “Safeguarding radio frequency spectrum”, 2025.

[BB23] Christopher M Bishop and Hugh Bishop, *Deep learning: Foundations and concepts*, Springer Nature, 2023.

[BBJ$^+$22] Alexandra Bringer, Annie Bete, Joel T. Johnson, Priscilla N. Mohammed, and Sidharth Misra, “Using machine learning algorithms to detect and classify RFI: Illustration using soil moisture active/passive (SMAP) data”, 2022.

[BJMP17] Alexandra Bringer, Joel T. Johnson, Priscilla N. Mohammed, and Jeffrey R. Piepmeier, “Performance of SMAP radiometer RFI detection algorithms and analysis of residual RFI sources”, in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 1243–1246, 2017.

[BS19] Jedda Boyle and Alessio Sclocco, “Edge-thresholding: A new thresholding based real-time RFI-mitigation algorithm for transient detection”, in *2019 RFI Workshop - Coexisting with Radio Frequency Interference (RFI)*, pp. 1–5, 2019.

[Cho20] François Chollet, “Variational autoencoder”, 2020.

[CM91] John C Curlander and Robert N McDonough, *Synthetic aperture radar*, Wiley, New York, 1991.

[CW05] Ian G Cumming and Frank H Wong, “Digital processing of synthetic aperture radar data”, *Artech house*, vol. 1, no. 3, pp. 108–110, 2005.

[dHHA19] P. de Haan, M. Helder, and J. Aalders, “Impact of Radio Frequency Interference on Radar and Radiometer Systems for Earth Observation Applications”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, pp. 1–15, March 2019.

[DHS22] Abdelkader Dairi, Fouzi Harrou, and Ying Sun, “Efficient driver drunk detection by sensors: A manifold learning-based anomaly detector”, *IEEE Access*, vol. 10, pp. 119001–119012, 2022.

[dMSLV16] Paolo de Matthaeis, Yan Soldo, and David M. Le Vine, “Analysis of RFI statistics for Aquarius RFI detection and mitigation improvements”, in *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 824–825, 2016.

[Ell20] J.R. Elliott, “Earth Observation for the Assessment of Earthquake Hazard, Risk and Disaster Management”, *Surveys in Geophysics*, vol. 41, no. 6, pp. 1323–1354, 11 2020.

[FBKO23] Chris Finlay, Bruce A Bassett, Martin Kunz, and Nadeem Oozeer, “Trajectory-based RFI subtraction and calibration for radio interferometry”, *Monthly Notices of the Royal Astronomical Society*, vol. 524, no. 3, pp. 3231–3251, 07 2023.

[Fin68] Johnson Robert S. Finn, Harry M., “Adaptive detection mode with threshold control as a function of spatially sampled-clutter-level estimates”, vol. 29, no. 3, pp. 414–463, 1968.

[FPR+21] Niccolò Franceschi, Riccardo Piantanida, Andrea Recchia, Davide Giudici, Andrea Monti-Guarnieri, Nuno Miranda, and Clément Albinet, “RFI monitoring and classification in c-band using sentinel-1 noise pulses”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 12, pp. 10425–10434, 2021. ARESYS, Italy; Politecnico di Milano, Italy; ESA/ESRIN, Italy.

[GAFC23] Søren G. Graabæk, Emil Vincent Ancker, Andreas Rune Fugl, and Anders Lyhne Christensen, “An experimental comparison of anomaly detection methods for collaborative robot manipulators”, *IEEE Access*, vol. 11, pp. 65834–65848, 2023.

[GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, 2016. `http://www.deeplearningbook.org`.

[GCW+15] Hugh Griffiths, Lawrence Cohen, Simon Watts, Eric Mokole, Chris Baker, Mike Wicks, and Shannon Blunt, “Radar spectrum engineering and management: Technical and regulatory issues”, *Proc. IEEE*, vol. 103, no. 1, pp. 85–102, 2015.

[Gil18] Enrique Bonet Gil, *10 - ASCE Test Cases*, Elsevier, 2018.

[GRS19] IEEE GRSS, “IEEE GRSS RFI observations display system”, 2019.

[HCRDR19] Stephen Harrison, Rory Coles, Timothy Robishaw, and David Del Rizzo, “RFI novelty detection using machine learning techniques”, in *2019 RFI Workshop - Coexisting with Radio Frequency Interference (RFI)*, pp. 1–6, 2019.

[HOS+20] H. Hébert, G. Occhipinti, F. Schindelé, A. Gailler, B. Pinel-Puysségur, H. K. Gupta, L. Rolland, P. Lognonné, F. Lavigne, E. Meilianda, S. Chapkanski, F. Crespon, A. Paris, P. Heinrich, A. Monnier, A. Jamelot, and D. Reymond, “Contributions of space missions to better tsunami science: Observations, models and warnings”, *Surveys in Geophysics*, vol. 41, pp. 1535–1581, 2020.

[IL19] Stephen Itschner and Xin Li, “Radio frequency interference (RFI) detection in instrumentation radar systems: a deep learning approach", in *2019 IEEE Radar Conference (RadarConf)*, pp. 1–5, 2019.

[JYB16] Ahsan Jalil, Hassan Yousaf, and Muhammad Iram Baig, "Analysis of cfar techniques", in *2016 13th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, pp. 654–659, 2016.

[KBC+14] Andreas Kääb, Tobias Bolch, Kimberly Casey, Torborg Heid, Jeffrey Kargel, Gregory Leonard, Frank Paul, and Bruce Raup, "Glacier mapping and monitoring using multispectral data", 07 2014.

[KFH+05] Gerhard Krieger, Hauke Fiedler, I. Hajnsek, Michael Eineder, Marian Werner, and Alberto Moreira, "TanDEM-X: Mission concept and performance analysis", volume 7, pp. 4890– 4893, 08 2005.

[KH10] Helmut Klausing and Wolfgang Holpp, *Radar mit realer und synthetischer Apertur: Konzeption und Realisierung*, Walter de Gruyter, 2010.

[KK24] Vladimír Kunc and Jiří Kléma, "Three decades of activations: A comprehensive survey of 400 activation functions for neural networks", 2024.

[KSSB19] Steen S. Kristensen, Niels Skou, Sten S. Søbjærg, and Jan E. Balling, "Developments of RFI detection algorithms and their application to future european spaceborne systems", in *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, pp. 4451–4454, 2019.

[KW19] Diederik P. Kingma and Max Welling, "An introduction to variational autoencoders", *CoRR*, vol. abs/1906.02691, 2019.

[LDO+19] Alvaro Llorente, Elena Daganzo, Roger Oliva, Ekhi Uranga, Yann Kerr, P. Richaume, Antonio de la Fuente, Manuel Martin-Neira, and S. Mecklenburg, "Lessons learnt from SMOS RFI activities after 10 years in orbit: RFI detection and reporting to claim protection and increase awareness of the interference problem in the 1400–1427 mhz passive band", in *2019 RFI Workshop - Coexisting with Radio Frequency Interference (RFI)*, pp. 1–6, 2019.

[LTZ09] Fei Tony Liu, Kai Ting, and Zhi-Hua Zhou, "Isolation forest", pp. 413 – 422, 01 2009.

[LWZ19] W. R. Lamb, M. S. Wolfe, and D. L. Zimovetsky, "Commercial cameras for satellite-based earth observation", *Remote Sensing of Environment*, vol. 237, pp. 111–124, 2019.

[LZL+22] Ning Li, Hengrui Zhang, Zongsen Lv, Lin Min, and Zhengwei Guo, "Simultaneous screening and detection of RFI from massive SAR images: A case study on european Sentinel-1", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–17, 2022.

[MACC24] Robert Muja, Andrei Anghel, Remus Cacoveanu, and Silviu Ciochina, "SAR interference mitigation using a filter bank time-frequency distribution and l-statistics", in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, pp. 7938–7942, 2024.

[Mor00] Alberto Moreira, *Radar mit synthetischer Apertur*, 2000.

[Oy22] HARP Technologies Oy, "Interference detection, classification and cancellation from space", 2022.

[PPQC24] Adrian Perez-Portero, Jorge Querol, and Adriano Camps, "Classification and characterization of RFI events for passive earth observation bands", in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, pp. 743–745, 2024.

[PQPC20] Adrian Pérez, Jorge Querol, Hyuk Park, and Adriano Camps, "Radio-frequency interference location, detection and classification using deep neural networks", in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, pp. 6977–6980, 2020.

[QCP$^+$21] Jorge Querol, Adriano Camps, Adrian Pérez, Roger Oliva, Raul Onrubia, Juan I. Ramirez-Martinez, Alberto Zurita, Martin Suess, and Manuel Martin-Neira, "A pre-correlation RFI mitigation algorithm for l-band interferometric radiometers", in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pp. 1819–1822, 2021.

[S$^+$18] Osvaldo Simeone et al., "A brief introduction to machine learning for engineers", *Foundations and Trends® in Signal Processing*, vol. 12, no. 3-4, pp. 200–431, 2018.

[SdMLV16] Yan Soldo, Paolo de Matthaeis, and David M. Le Vine, "L-band RFI in Japan", in *2016 Radio Frequency Interference (RFI)*, pp. 111–114, 2016.

[SGS18] K. Scharf, J. Gaultier, and B. Stenberg, "Comparing the Detection Capabilities of Radar and Radiometer Systems for Land Surface Monitoring", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, pp. 1–13, March 2018.

[SKAH$^+$24] Maximilian Schandri, Thomas Kraus, Akram Al-Hourani, Nermine Hendy, Ferdi Ganda Kurnia, and Markus Bachmann, "Analysis of radio frequency interferences in TerraSAR-X products", in *EUSAR 2024; 15th European Conference on Synthetic Aperture Radar*, pp. 916–919, 2024.

[SLVdM19] Yan Soldo, David Le Vine, and Paolo de Matthaeis, “An approach to address residual “hot spots” in SMAP RFI-filtered data”, in *2019 RFI Workshop - Coexisting with Radio Frequency Interference (RFI)*, pp. 1–3, 2019.

[STG21] Siddhartha Subray, Stefan Tschimben, and Kevin Gifford, “Towards enhancing spectrum sensing: Signal classification using autoencoders”, *IEEE Access*, vol. 9, pp. 82288–82299, 2021.

[TSHW19] Mingliang Tao, Jia Su, Yan Huang, and Ling Wang, “Mitigation of radio frequency interference in synthetic aperture radar data: Current status and future trends”, *Remote Sensing*, vol. 11, no. 20, 2019.

[Uni09] International Telecommunication Union, “Itu-r 2009”, 2009.

[Uni25] ITU International Telecommunication Union, “Space connect: The rise of leo satellite constellations”, 2025.

[VDG+19] Aravind Venkitasubramony, Eryan Dai, Albin Gasiewski, Maciej Stachura, and Jack Elston, “RFI detection and mitigation in an suas based l band correlation radiometer for soil moisture measurements”, in *2019 RFI Workshop - Coexisting with Radio Frequency Interference (RFI)*, pp. 1–4, 2019.

[Wan21] Jianping Wang, “Cfar-based interference mitigation for fmcw automotive radar systems”, 2021.

[WHG+25] Tingting Wei, Xingwang Hu, Zhengwei Guo, Gaofeng Shu, Yabo Huang, and Ning Li, “A two-stage method for screening pulse RFI in SAR raw data alternating the use of time and frequency domains”, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 4331–4346, 2025.

[WWF+19] Xinxin Wang, Xiang Wang, Jianchao Fan, Jianhua Zhao, Yu Wang, and Enbo Wei, “Automatic detection and identification of RFI sources for SMAP satellite polarized data based on IDL”, in *2019 Tenth International Conference on Intelligent Control and Information Processing (ICICIP)*, pp. 76–80, 2019.

[YLL+24] Huizhang Yang, Ping Lang, Xingyu Lu, Shengyao Chen, Feng Xi, Zhong Liu, and Jian Yang, “Robust block subspace filtering for efficient removal of radio interference in synthetic aperture radar images”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–12, 2024.

[ZLX+24] Jiayi Zhao, Guisheng Liao, Hang Xing, Boyang Jia, Sining Liu, Yongliang Wang, and Xuepan Zhang, “Range-frequency domain iterative notch filtering method for RFI mitigation in SAR”, *IEEE Geoscience and Remote Sensing Letters*, vol. 21, pp. 1–5, 2024.

Karlsruhe Institute of Technology | Deutsches Zentrum für Luft- und Raumfahrt

B.Eng. Caterina Simone Maria Klumpp
Master Thesis: Analysis and Mitigation of Radio Frequency Interferences in X-band SAR Data
Zeitraum: 11.12.2024 – 02.05.2025
Hauptreferent: Prof. Dr.-Ing. Thomas Zwick
Betreuer: Dr.-Ing. Markus Bachmann , M.Sc. Thomas Kraus, M.Sc. Maximilian Schandri

Abstract: In recent years, terrestrial RFI from sources such as airports, harbors or air surveillance systems, as well as spaceborne RFIs from growing multisatellite constellations have increasingly impacted TerraSAR-X data.

This thesis investigates methods to detect and mitigate these interferences in raw data. Therefore, machine learning methods, including iForest, autoencoder, and VAE architectures, are evaluated and compared regarding their detection performance. A clearly superior detection performance of the autoencoder for this particular case is found.

The performance of conventional SIR, CA-CFAR filtering and the autoencoder are further assessed and compared in terms of detection performance and mitigation effectiveness. As a reference a conventional SIR filter of $18\,\mathrm{dB}$ is determined for the scenes at hand and set as the benchmark. In case of NBI the CA-CFAR detects RFIs in additional $0.6\,\%$ of the image's samples and the autoencoder an additional $1.3\,\%$. In case of WBI the CA-CFAR detects $7.93\,\%$ less than the SIR filter and the autoencoder an additional $2.7\,\%$. Therefore, the detection rate of the autoencoder not only outperforms the conventional methods but it is also more robust. However, blanking the isolated RFIs in each rangeline's PSD for mitigation shows that the improved detection rate of the autoencoder to the SIR benchmark visually has a neglectable impact.

Finally, a data set derived from gloablly distributed SAR acquisitions is analyzed based on occupied bandwidth, RFI power, and the number of RFIs in its PSD. This enables a global mapping of RFI occurrences in TerraSAR-X data from 2023 and 2024, along with an initial analysis of the nature of these RFIs and hence proves the power of the developed autoencoder.

Postal address:	Institute of Radio Frequency Engineering and Electronics Kaiserstraße 12 D-76131 Karlsruhe	Tel.: Secy.: Fax: Email: Web:	+49 (0) 721 608 4 2522 +49 (0) 721 608 4 2523 +49 (0) 721 608 4 5027 info@ihe.kit.edu www.ihe.kit.edu
Building:	Engesserstraße 5, Geb. 30.10		