Restoring dynamic spectra corrupted by radio frequency interference : Application to pulsar observation by SKA-pathfinder low-frequency NenuFAR

Xiao Zhang

To cite this version:

Xiao Zhang. Restoring dynamic spectra corrupted by radio frequency interference : Application to pulsar observation by SKA-pathfinder low-frequency NenuFAR. Astrophysics [astro-ph]. Université d'Orléans, 2024. English. NNT : 2024ORLE1027 . tel-04839556

HAL Id: tel-04839556
https://theses.hal.science/tel-04839556v1
Submitted on 16 Dec 2024

UNIVERSITÉ D'ORLÉANS

ÉCOLE DOCTORALE ÉNERGIE MATÉRIAUX SCIENCES
DE LA TERRE ET DE L'UNIVERS

Laboratoire de Physique et de Chimie
de l'Environnement et de l'Espace

THÈSE présentée par :

Xiao ZHANG

soutenue le : 05 Novembre 2024

pour obtenir le grade de : Docteur de l'Université d'Orléans

Discipline/ Spécialité : Sciences de l'Univers

Restoring dynamic spectra corrupted by radio frequency interference

–

Application to pulsar observation by SKA-pathfinder low-frequency NenuFAR

THÈSE dirigée par :

M. COGNARD Ismaël	Directeur de recherche, Université d'Orléans
M. DOBIGEON Nicolas	Professeur, Toulouse INP

RAPPORTEURS :

M. EL KORSO Mohammed Nabil	Professeur, Université Paris-Saclay
M. HESSELS Jason	Professor, University of Amsterdam

JURY :

M. DUDOK DE WIT Thierry	Professeur, Université d'Orléans, Président du jury
M. EL KORSO Mohammed Nabil	Professeur, Université Paris-Saclay
M. HESSELS Jason	Professor, University of Amsterdam
M. GRATADOUR Damien	Chargé de recherche, Observatoire de Paris
M. COGNARD Ismaël	Directeur de recherche, Université d'Orléans
M. DOBIGEON Nicolas	Professeur, Toulouse INP

Remerciements

Je voudrais exprimer ma profonde gratitude à tous ceux qui m'ont aidé au cours de ma thèse pendant ces trois dernières années.

Tout d'abord, je veux dire merci à mon directeur de thèse principal, Nicolas, pour son soutien indéfectible, sa patience et ses précieux conseils au cours de ma thèse. Il m'a non seulement enseigné la manière de penser face aux problèmes, mais il a également partagé de manière désintéressée son expérience inestimable en traitement du signal, ce qui a eu un impact profond sur moi. Je veux également remercier mon co-directeur de thèse, Ismaël, pour la patience dont il a fait preuve en m'expliquant ses connaissances sur les pulsars et NenuFAR.

Ensuite, je voudrais remercier tous les membres de l'équipe pour toute l'aide qu'ils m'ont apportée dans mon travail et dans ma vie.

Enfin, je voudrais remercier mes parents pour leur soutien constant, pour m'avoir soutenu dans mes études en France voici dix ans et pour m'avoir encouragé à poursuivre mes études de thèse voilà trois ans.

Contents

Liste des figures iv

Liste des tables vi

Introduction 1

I Background 5

1 Pulsar 7
- 1.1 A star made of neutrons 7
- 1.2 Discovery of pulsar 8
- 1.3 Characteristics of pulsars 9
  - 1.3.1 Lighthouse effect of pulsars 9
  - 1.3.2 Pulsar rotation and emission evolution 10
  - 1.3.3 Radio beam emission 11
  - 1.3.4 Interstellar dispersion 13
  - 1.3.5 Population of pulsars 14

2 Radio astronomy telescope – NenuFAR 17
- 2.1 Radio astronomy and telescopic techniques 17
- 2.2 NenuFAR 18
  - 2.2.1 NenuFAR's mini-array and receiver 18
- 2.3 NenuFAR in pulsar observation 19

3 Radio Frequency Interference (RFI) 23
- 3.1 RFI 23
- 3.2 Importance and challenges of RFI mitigation 24
- 3.3 RFI mitigation at various observation stages 24
  - 3.3.1 Observatory prevention 24
  - 3.3.2 Pre-detection strategies 25
  - 3.3.3 Pre-correlation techniques 26
  - 3.3.4 Post-correlation methods 26
- 3.4 Overview of post-correlation RFI mitigation methods 27
  - 3.4.1 Traditional RFI mitigation techniques 27
  - 3.4.2 Data-driven RFI mitigation approaches 28
  - 3.4.3 Summary and challenges 30

II Contribution to RFI mitigation 33

4 Framework for simulating RFI corrupted dynamic spectra observation 35
- 4.1 Introduction and motivation 35
- 4.2 Observation model 35
- 4.3 Simulation of pulsar signal 36

4.3.1 Template-Based Model Enhancement . . . 38
4.3.2 Dispersion measurement simulation . . . 40
4.4 Simulation of RFI signal . . . 40
4.4.1 Extended RFI in time and frequency . . . 41
4.4.2 Narrow-band transient RFI . . . 41
4.5 Generation of simulated data sets . . . 42
4.5.1 Simulated data generation process . . . 42
4.5.2 Simulation parameter configuration . . . 42
4.5.3 Simulation scenario construction . . . 44

5 Introduction to deep learning and convolutional neural networks 47
5.1 Introduction to deep learning . . . 47
5.2 Foundations of neural networks . . . 47
5.2.1 Perceptron . . . 48
5.2.2 Multi-layer perceptrons . . . 49
5.2.3 Activation function . . . 50
5.2.4 Loss function . . . 51
5.2.5 Back propagation . . . 52
5.2.6 Optimization . . . 53
5.3 Convolutional neural networks . . . 54
5.3.1 Convolutional layers . . . 55
5.3.2 Pooling layer . . . 56
5.3.3 Applications of CNN in vision tasks . . . 56

6 Deep learning for RFI mitigation 63
6.1 Formulating RFI mitigation as image segmentation . . . 63
6.2 Formulating RFI mitigation as image restoration . . . 65
6.3 Proposed approach: RFI-DRUNet . . . 65
6.3.1 Architecture of RFI-DRUNet . . . 66
6.3.2 Core building block of the RFI-DRUNet architecture . . . 67

7 Experimental framework and results 71
7.1 Experimental framework . . . 71
7.1.1 Implementation details . . . 72
7.1.2 Quantitative metrics . . . 72
7.2 Experimental results . . . 73
7.2.1 Restoration results . . . 73
7.2.2 Detection results . . . 81
7.2.3 Ablation study . . . 82
7.2.4 Illustration on a real observation . . . 84
7.2.5 Discussion . . . 85

8 Application to pulsar TOA estimation 89
8.1 Numerical experiments . . . 91

Conclusion 95

Bibliographie 99

List of Figures

1.1 Chart of the first recorded pulsar (`PSR B1919+21`). 8
1.2 Toy model of a rotating pulsar. 10
1.3 Representation the geometry of pulsar emission beams. 12
1.4 Illustration the dispersion effect on pulsar signals as they propagate through the interstellar medium. 13
1.5 The diagram $P - \dot{P}$ of pulsar. 14
1.6 Schematic representation of the evolution of pulsar binaries. 15

2.1 Spectral window of radio waves detectable from Earth. 17
2.2 Schematic distribution diagram of NenuFAR's mini-arrays. 19
2.3 Schematic diagram of NenuFAR receiver. 20

3.1 Examples of typical RFI waveforms in time domain. 23
3.2 Map of protection zones for the Nançay radio observatory. 25

4.1 An illustration of the dynamic spectrum generated on the basis of the proposed model delineated in Equation (4.1). 36
4.2 Illustration of pulse profile evolution with Frequency for two pulsar signal. 38
4.3 Pulse profiles of the pulsar B1919+21, which were recorded over a range of frequencies by NenuFAR telescope. 39
4.4 Examples of the dynamic spectrum generated according to the proposed protocol for the 4 cases of S_2. 45

5.1 Diagram representation of the perceptron. 48
5.2 Diagram representing the flow of data through a multi-layer perceptron. 49
5.3 Diagram representing the computational process of a single MLP layer with two neurons . 50
5.4 An illustration of commonly used activation functions. 51
5.5 Network Architecture and Data Flow of LeNet. 55
5.6 Demonstration of the convolution operation in a CNN. 56
5.7 Demonstration of the max pooling operation in a CNN. 57
5.8 Simplified illustration of a convolutional neural network workflow for image classification. 58
5.9 Illustration of a CNN with an encoder-decoder structure workflow for image segmentation. 59

6.1 Illustration of U-Net applied for RFI mitigation. 64
6.2 Architecture of the proposed RFI-DRUNet network. 66
6.3 Illustration of residual connection a CNN. 67
6.4 Demonstration of the transposed convolution operation in a CNN. 68

7.1 Visual comparison of spectra restored by $\mathsf{S}_1\mathsf{C}_\mathrm{A}$ across four test data sets in scenario S_1. . 75
7.2 Visual comparison of RFI-corrupted spectrum in $\mathsf{S}_1\mathsf{C}_\mathrm{D}$ and its restored spectra provided by models $\mathsf{S}_1\mathsf{C}_\mathrm{B}$ to $\mathsf{S}_1\mathsf{C}_\mathrm{D}$. 76
7.3 Visual comparison of spectra restored by $\mathsf{S}_2\mathsf{C}_\mathrm{A}$ across four test data sets in scenario S_2. . 77
7.4 Visual comparison of RFI-corrupted spectrum in $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ and its restored spectra provided by models $\mathsf{S}_2\mathsf{C}_\mathrm{B}$ to $\mathsf{S}_2\mathsf{C}_\mathrm{D}$. 77

7.5 Visual comparison of RFI-corrupted spectrum in $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ and its restored spectra provided by model $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ and Oracle detector. 78
7.6 Illustration of restoration performance (in terms of PSNR) in validation. 79
7.7 Visual comparison of spectra with sinusoidal RFI and their restored spectra provided by $\mathsf{S}_2\mathsf{C}_\mathrm{D}$. 80
7.8 Visual comparison of the results for simulated spectrum from $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ provided by the compared methods. 82
7.9 Illustrating the trade-off between model's complexity (FLOPs) and restoration performance. 84
7.10 Results of visual comparison with RFI-DRUNet and U-Net on a real observation. 86

8.1 Scatter plot of the TOA estimation uncertainty versus SNR. 91
8.2 TOA estimation errors in terms of MSE (top) and MAR (bottom) as function of the uncertainty threshold. 92

List of Tables

4.1 Notations used to describe the RFI-corrupted dynamic spectra generated by the simulation protocol. 37
4.2 Parameters of the simulation associated with the pulsar signal. 43
4.3 Parameters of the simulation associated with the RFI signal. 44
4.4 Size of the generated training, validation and testing sets. 44

6.1 Details of the layers of the proposed RFI-DRUNet network. 67

7.1 Confusion matrix for RFI detection. 72
7.2 Restoration performance in Scenario 1. 74
7.3 Restoration performance in Scenario 2. 76
7.4 Restoration performance in Out-of-distribution data sets for different noise level. 80
7.5 Restoration performance in Out-of-distribution data sets for different RFI profiles. 81
7.6 Detection performance of compared algorithms. 82
7.7 Ablation study results. 83
7.8 Restoration performance of compared algorithms. 87
7.9 Number of parameters of the compared models. 87

Introduction

Since their discovery in 1967 (Hewish et al., 1968), pulsars, which are rapidly spinning neutron stars, have been a focus of astronomical research. These celestial objects, with their highly stable rotational periods, serve as natural laboratories for testing fundamental physics and as probes for gravitational wave detection. The stability of their periodic signals makes them invaluable tools for various astrophysical studies, particularly in the context of gravitational wave astronomy. The recent advancements in low-frequency radio astronomy have opened up a new chapter in the observation and understanding of pulsars. With start of the art instrumentations installed on mid-frequency radio telescopes, the detection of low-frequency gravitational waves is becoming an increasingly realistic goal.

The low-frequency domain, specifically below 100 MHz, is vital for pulsar observations thanks to its ability to provide insights into the scattering and dispersion effects in the interstellar medium. NenuFAR (Zarka et al., 2018), as a low-frequency pathfinder to the Square Kilometre Array (SKA) (Dewdney et al., 2009), plays a pivotal role in pulsar observations (Bondonneau et al., 2021). It provides a specialized observatory platform, thereby enhancing our ability to explore this less-charted frequency range below 100 MHz.

While NenuFAR displays considerable prospective, and the utilisation of low-frequency observations offers a promising potential, the presence of radio frequency interference (RFI) poses a significant challenge. RFI originating from various human-made sources, including telecommunications and broadcasting, can corrupt the faint signals from pulsars, thereby complicating the analysis and potentially leading to erroneous scientific conclusions. The threat of RFI to pulsar observations is particularly pronounced in the low-frequency range, where the interference often dominates over the astronomical signals.

To tackle this problem, traditional RFI mitigation methods, which encompass both hardware and software techniques, have been deployed. These conventional approaches frequently necessitate a substantial degree of prior knowledge regarding the signal characteristics, and they often demonstrate limited generalisation capabilities, particularly in complex and variable RFI environments. Their efficacy may be limited by the necessity for manual parameter tuning and their ability to accurately identify RFI in diverse scenarios.

Deep learning has introduced a new framework for RFI mitigation, with convolutional neural networks (CNN) successfully deployed for detecting and mitigating RFI in radio astronomical data. Nevertheless, the effective implementation of these techniques is confronted with two considerable obstacles.

Firstly, the lack of labelled training data is a major obstacle to the application of supervised deep learning approaches. Although there exist simulated data sets and some manually labelled datasets, these are often not specific to pulsar observation, which limits the training potential for RFI mitigation in this context. Secondly, a common challenge faced by both traditional and current deep learning based solution methods is the identification and removal of RFI-corrupted data segments. While this strategy is effective in identifying RFI, it fails to restore the corrupted signals, thereby compromising the integrity and completeness of the observational data.

In response to these challenges, this thesis introduces a novel framework to address the lack of labeled data. We propose a simulation framework designed to generate realistic, RFI-corrupted pulsar observation data, complete with clean pulsar signals. This framework enables the training of supervised deep learning models, thus enhancing the applicability of these methods to RFI mitigation in pulsar observations.

Furthermore, we propose a paradigm shift in RFI mitigation , treating it as an image restoration task rather than detection and removal. By formulating the problem in this manner, we can develop deep learning models that not only identify RFI but also restore the corrupted signals, preserving the integrity of the astronomical data. Our proposed approach leverages the power of deep denoising networks, demonstrating their potential for effective RFI mitigation and signal restoration.

The contributions of this work are manifold. Firstly, we combine the tasks of detection and restoration by employing a denoising problem network, making it a reality. Secondly, we introduce a simulation framework that enables the use of supervised deep neural networks for RFI mitigation. Thirdly, the proposed method, based on a deep denoising network, achieves remarkable results in signal restoration while maintaining competitive RFI detection capabilities. Lastly, we demonstrate the utility of our approach by improving the estimation of pulsar time of arrival (TOA), showing that the restored signals closely match the performance of RFI-free signals.

The subsequent chapters of this thesis are structured as follows: Chapter 1 provides a comprehensive background on pulsars, emphasizing their significance in contemporary astronomy and their role in gravitational wave detection. Chapter 2 examines the technical specifications and observational capabilities of NenuFAR, highlighting its unique advantages for low-frequency pulsar observations. Chapter 3 reviews the existing methods for RFI mitigation, discussing their limitations and the rationale behind the need for novel approaches. Chapter 4 introduces the proposed simulation framework for generating RFI-corrupted pulsar data. Chapter 5 provides an in-depth review of deep learning and convolutional neural networks, setting the stage for the presentation of our proposed method. Chapter 6 details the proposed RFI-DRUNet network and its application to RFI mitigation. Finally, Chapter 7 presents the experimental framework, results, and discussions, demonstrating the effectiveness of our approach in restoring RFI-corrupted signals and improving pulsar TOA estimation.

This thesis aims to contribute to the field of radio astronomy by addressing the pressing issue of RFI in pulsar observations, offering a novel perspective on RFI mitigation, and enhancing the potential for gravitational wave detection through improved pulsar timing.

Introduction (français)

Depuis leur découverte en 1967 Hewish et al. (1968), les pulsars, ces étoiles à neutrons en rotation rapide, sont au cœur des recherches astronomiques. Leur période de rotation extrêmement stable en fait des laboratoires naturels pour tester la physique fondamentale et des outils précieux pour la détection des ondes gravitationnelles. La stabilité de leurs signaux périodiques les rend essentiels pour mener des études astrophysiques variées, notamment dans le domaine de l'astronomie des ondes gravitationnelles. Les récentes avancées en radioastronomie à basse fréquence ont ouvert un nouveau chapitre dans l'observation et la compréhension des pulsars. Avec des intrumentations à la pointe installées sur les radiotélescopes décimétriques, la détection desondes gravitationnelles à basse fréquence devient de plus en plus réalisable.

Le domaine des basses fréquences, en dessous de 100 MHz, est crucial pour les observations de pulsars car il offre des informations sur les effets de diffusion et de dispersion dans le milieu interstellaire. NenuFAR (Zarka et al., 2018), en tant que précurseur à basse fréquence pour le projet de radiotélescope *Square Kilometre Array* (SKA) (Dewdney et al., 2009), joue un rôle clé dans ces observations en offrant une plateforme spécialisée qui améliore notre capacité à explorer cette gamme de fréquences encore peu étudiée.

Cependant, la présence d'interférences radioélectriques (RFI) constitue un défi majeur. Ces interférences, provenant de diverses sources humaines telles que les télécommunications et la radiodiffusion, peuvent altérer les faibles signaux des pulsars, compliquant ainsi leur analyse et risquant de fausser les conclusions scientifiques. Les RFI sont particulièrement problématiques dans la gamme des basses fréquences où elles dominent souvent les signaux astronomiques.

Pour faire face à ce défi, des méthodes traditionnelles d'atténuation des RFI, combinant techniques matérielles et logicielles, ont été mises en œuvre. Ces méthodes nécessitent souvent une connaissance préalable approfondie des signaux et présentent des capacités limitées de généralisation, surtout dans des environnements de RFI complexes. Leur efficacité est parfois entravée par la nécessité d'ajuster manuellement les paramètres et par leur incapacité à identifier avec précision les RFI dans des scénarios variés.

L'apprentissage profond a introduit un nouveau cadre pour le traitement des RFI, avec des réseaux de neurones convolutifs réussissant à détecter et atténuer ces interférences dans les données radioastronomiques. Toutefois, l'application efficace de ces techniques se heurte à deux obstacles principaux. D'abord, le manque de données d'apprentissage étiquetées est un obstacle majeur pour les approches supervisées. Les ensembles de données disponibles sont souvent simulés ou manuellement étiquetés, et ne sont pas toujours spécifiques à l'observation des pulsars. Ensuite, que ce soit pour les méthodes traditionnelles ou celles basées sur l'apprentissage profond, la suppression des données corrompus par les RFI reste un défi, car ces techniques ne restaurent pas les signaux corrompus, compromettant ainsi l'intégrité des données d'observation.

Face à ces défis, cette thèse propose un cadre innovant pour résoudre le problème du manque de données étiquetées. Nous introduisons un cadre de simulation permettant de générer des données d'observation de pulsars réalistes, corrompues par des RFI, tout en conservant des signaux de pulsars purs. Ce cadre facilite la formation de modèles d'apprentissage profond supervisé, améliorant ainsi l'applicabilité des méthodes pour la mitigation des RFI dans le contexte des pulsars.

En outre, nous proposons de traiter la mitigation des RFI comme une tâche de restauration d'image plutôt que de simple détection et suppression. En formulant le problème de cette manière, nous développons des modèles d'apprentissage profond capables non seulement d'identifier les RFI, mais aussi de restaurer les signaux corrompus, préservant ainsi l'intégrité des données astronomiques. Nous exploitons le potentiel des réseaux de dénoising profonds pour démontrer l'efficacité de notre approche dans la mitigation des RFI et la restauration des signaux.

Les contributions de ce travail sont multiples. Premièrement, nous combinons détection et restauration en utilisant un réseau de nettoyage de problèmes. Deuxièmement, nous introduisons un cadre de simulation permettant l'utilisation de réseaux neuronaux profonds supervisés pour la mitigation des RFI. Troisièmement, la méthode basée sur un réseau de dénoising profond montre des résultats remarquables en matière de restauration des signaux tout en maintenant une détection des RFI efficace. Enfin, nous démontrons l'amélioration de l'estimation des temps d'arrivée des pulsars, montrant que les signaux restaurés se rapprochent des performances des signaux sans RFI.

Les chapitres suivants sont structurés comme suit. Le Chapitre 1 fournit un contexte complet sur les pulsars, soulignant leur importance dans l'astronomie contemporaine et leur rôle dans la détection des ondes gravitationnelles. Le Chapitre 2 examine les spécifications techniques et les capacités d'observation de NenuFAR, mettant en avant ses avantages uniques pour les observations de pulsars à basse fréquence. Le Chapitre 3 passe en revue les méthodes existantes d'atténuation des RFI, en discutant de leurs limitations et de la nécessité de nouvelles approches. Le Chapitre 4 introduit le cadre de simulation proposé pour générer des données de pulsars corrompues par des RFI. Le Chapitre 5 fournit une revue approfondie de l'apprentissage profond et des réseaux de neurones convolutifs. Le Chapitre 6 détaille le réseau RFI-DRUNet et son application à l'atténuation des RFI. Enfin, le Chapitre 7 présente le cadre expérimental, les résultats et les discussions, démontrant l'efficacité de notre approche pour restaurer les signaux corrompus par les RFI et améliorer l'estimation des temps d'arrivée des pulsars.

Cette thèse vise à contribuer à la radioastronomie en abordant la question cruciale des RFI dans les observations de pulsars, en offrant une nouvelle perspective pour leur atténation et en renforçant le potentiel de détection des ondes gravitationnelles grâce à une meilleure chronométrie des pulsars.

Part I

Background

Chapter 1

Pulsar

1.1 A star made of neutrons

At the beginning of the 20th century, the quest to understand the composition of the atomic nucleus was a major focus of physics. Ernest Rutherford's experimental investigations identified the proton and formulated the proton-electron model of the atomic nucleus (Rutherford, 1911). This model suggested that the atomic nucleus was made up of protons and electrons. However, the observed mass of the nucleus far exceeded the combined mass of protons and electrons, suggesting the presence of additional, as yet undiscovered, particles within the nucleus. In 1932, Chadwick (1932) made a landmark discovery when he detected the neutron, a previously unknown particle characterised by its lack of charge and high penetrative power. This discovery was made possible by the experiment of bombarding beryllium with alpha particles. [1]

The detection of neutrons has facilitated the acquisition of novel insights into the physical processes occurring within stellar interiors. The concept of a neutron star was first introduced by Baade and Zwicky (1934), who formulated the hypothesis while examining the potential remnants of dense material following a supernova explosion. They postulated that the core of a star might undergo gravitational collapse following a supernova, resulting in a compact object predominantly constituted by neutrons, an entity now recognised as a neutron star.

Subsequently, in 1938, Lev Landau conducted research into the physical characteristics of neutron stars. His research included an investigation of their theoretical stability and probable composition, thereby establishing a crucial theoretical framework for the subsequent study of neutron stars.

Building upon these theoretical advancements, Oppenheimer and Volkoff (1939) employed the principles of general relativity and quantum mechanics in 1939 to analyse the equilibrium between gravitational and neutron degeneracy pressures within the context of neutron stars. They constructed a model of the internal structure of neutron stars and derived the theoretical upper boundary for the mass of a neutron star, a threshold subsequently designated as the Tolman-Oppenheimer-Volkov limit (TOV limit).

The study of neutron stars experienced a period of stagnation until 1967, when Pacini (1967) put forth the hypothesis that a rapidly spinning, highly magnetised, and extremely dense neutron star could emit electromagnetic radiation. This theoretical work provided an explanation for the previously observed extraordinarily luminous radio source in the Crab Nebula, which had been detected by Hewish and Okoye (1965). Subsequently, Hewish et al. (1968) identified a group of pulsars, building upon the theoretical insights that had been previously established.

At the time of their discovery, two principal theoretical models were proposed to account for the nature of pulsars. The initial theoretical model, proposed by Meltzer and Thorne (1966), postulated that radial oscillations in white dwarfs could generate periodic signals of several seconds' duration. In contrast, the second model, proposed by Pacini (1967), suggests that pulsars are rapidly rotating neutron stars. This hypothesis was subsequently refined by Gold (1968), who proposed that neutron stars, acting as rotating magnetic dipoles, emit relativistic particles and electromagnetic radiation, thereby causing an energy loss and subsequent deceleration of the star's rotation.

The detection of a pulsar (`PSR B0531+21`) with a period of 33 milliseconds in the Crab Nebula by

---

[1] This chapter has been partially inspired by Lorimer and Kramer (2004).

Staelin and Reifenstein III (1968) using the Green Bank Telescope constituted the inaugural empirical evidence in support of the neutron star hypothesis. Moreover, Richards and Comella (1969) employed the Arecibo Observatory to conduct observations that revealed details about the period of the Crab Pulsar, thereby substantiating the neutron star theory and advancing our understanding of these enigmatic celestial objects.

1.2 Discovery of pulsar

A pulsar is an enigmatic relic of a supernova explosion, representing a highly magnetized, rapidly rotating neutron star. These celestial objects emit beams of electromagnetic radiation from their magnetic poles, which, due to the star's rotation, sweep across the sky like a lighthouse beacon. The precision and regularity of these emissions are such that pulsars are often referred to as the most accurate natural clocks in the universe. The first pulsar was discovered in 1967 at the Mullard Radio Astronomy Observatory (MRAO) at the Cambridge University by Hewish et al. (1968). They parsed the signals into a series of pulses spaced 1.37 seconds apart by means of a fast strip chart recorder (in the Figure 1.1). Initially, these signals were thought to be man-made interference because they appeared to be unusually regular. However, through continued observations (Pilkington et al., 1968), astronomers have found that these signals have a fixed location in the sky and exhibit a very high degree of periodicity and persistence. These properties are not consistent with the common characteristics of terrestrial or man-made RFI, which usually vary over time or are observed in different locations. These observations therefore ruled out the possibility that the signals were RFI and led scientists to determine that they originated from some kind of astronomical object.

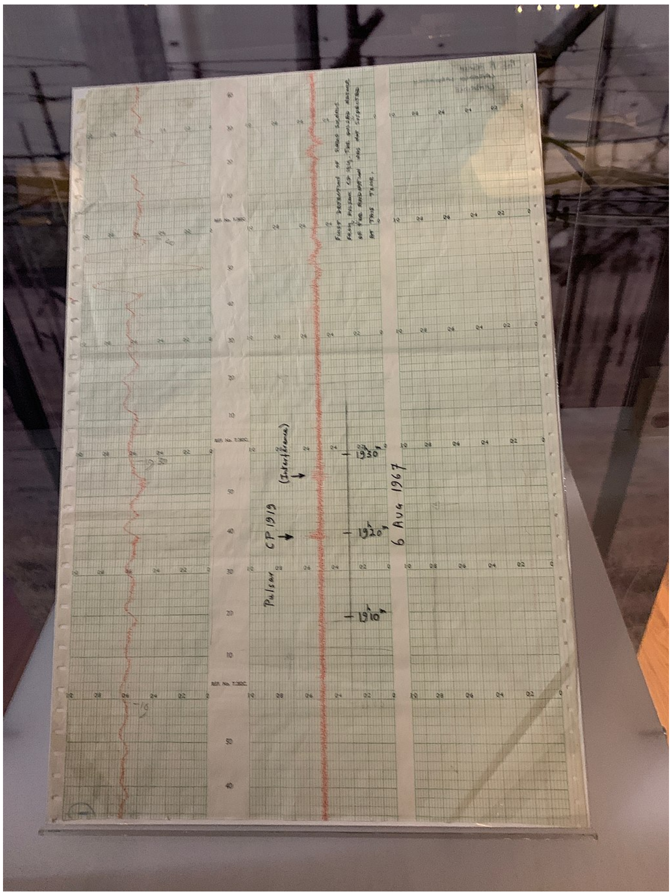

Figure 1.1: Chart where Jocelyn Bell Burnell first identified evidence of `PSR B1919+21`, now displayed at Cambridge University Library (picture extracted from Wikipedia[2]).

---

[2] By Billthom - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=77181280

With further study (Gold, 1968; Goldreich and Julian, 1969; Ostriker and Gunn, 1969; Gunn and Ostriker, 1969; Vila, 1969), the source of these regular pulses was soon deduced to be the neutron stars being sought, the collapsed cores of massive stars after supernova explosions. Such objects have powerful magnetic fields that produce highly collimated beams of electromagnetic radiation. The term "pulsar" is derived from "pulsating star", and is used to describe these regularly flashing objects. This discovery not only confirms the existence of neutron stars, but also has far-reaching implications for the fields of astronomy and physics, opening up new avenues for the study of extreme astrophysical conditions.

Hulse and Taylor (1975) discovered a binary system of radio pulsars now called `PSR B1913+16`, a discovery of great importance to astrophysics and gravitational wave astronomy. Through detailed observations of this binary star system, they found that the orbital periods of the pulsars are gradually decreasing, a phenomenon that is consistent with Einstein's general relativity predictions about gravitational wave radiation (Taylor et al., 1979; Taylor and Weisberg, 1982). Their study suggests that the two pulsars in the binary system are losing energy due to the emission of gravitational waves, causing them to move closer to each other, which triggers a reduction in the orbital period. This finding provides indirect evidence for the existence of gravitational waves. The exceptional rotational stability of pulsars makes them ideal astrophysical laboratories for conducting high-precision timing measurements. Studies have shown that pulsars can be utilized to accurately measure orbital parameters, providing a valuable tool for investigating the effects of gravitational waves (Gullahorn and Rankin, 1978; Weisberg et al., 2010). The discovery of the `PSR B1913+16` binary system is a milestone in the history of astrophysics, which not only confirms the existence of gravitational waves, but also lays a theoretical foundation for future gravitational wave detection and research. Note that, discovered in 2003 (Burgay et al., 2003), the double pulsar `J0737-3039A` is now the best system to test the Einstein's general relativity (Kramer et al., 2021).

In 1982, Don Backer and his team discovered `PSR B1937+21` (Backer et al., 1982), a pulsar with an exceptionally fast rotation period of only 1.6 milliseconds (MSPs). This fast rotation is currently explained, in an evolved binary system, by accretion of matter transferring orbital momentum to the neutron star, spinning it up to millisecond periods. The high rotational stability of millisecond pulsars makes them ideal astrophysical clocks for precise timing measurements (Davis et al., 1985; Taylor, 1991; Matsakis et al., 1997; Hobbs et al., 2012). In addition, they are being considered for use in navigation systems for deep space probes (Buist et al., 2011; Shuai, 2021). More than 90% of millisecond pulsars are found in binary systems, and their interactions with the companion star can lead to changes in the orbital parameters of pulsars, which are well suited for detecting gravitational waves (Manchester et al., 2005; Kramer et al., 2006; Arzoumanian et al., 2018; Reardon et al., 2021; Agazie et al., 2023; Antoniadis et al., 2023; Reardon et al., 2023).

1.3 Characteristics of pulsars

1.3.1 Lighthouse effect of pulsars

The lighthouse effect of pulsars is a vivid analogy that describes how these rapidly rotating neutron stars sweep their beams of radiation into space. This phenomenon can be intuitively understood through the "toy model" depicted in Figure 1.2 which illustrates the relationship between the magnetic poles of the pulsar, the radiation beams, and the observer.

Pulsars, possessing some of the most intense magnetic fields in the universe, exhibit field strengths ranging between 10^8 and 10^{13} Gauss. Such magnetic fields generate open magnetic field lines in the polar regions of the pulsar, in contrast to closed field lines that are entirely confined within the light cylinder. The light cylinder, a conceptual cylindrical boundary, is determined by the pulsar's rotation period with the radius calculated using the formula:

$$R_{\mathrm{LC}} = \frac{c}{\Omega} = \frac{c}{2\pi} P \tag{1.1}$$

where c represents the speed of light and P denotes the rotation period of the pulsar.

Within these open magnetic field line regions, particularly near the magnetic poles, charged particles (predominantly electrons and positrons) are accelerated to velocities approaching the speed of light, thereby acquiring high energies. These energetic particles subsequently emit photons through various

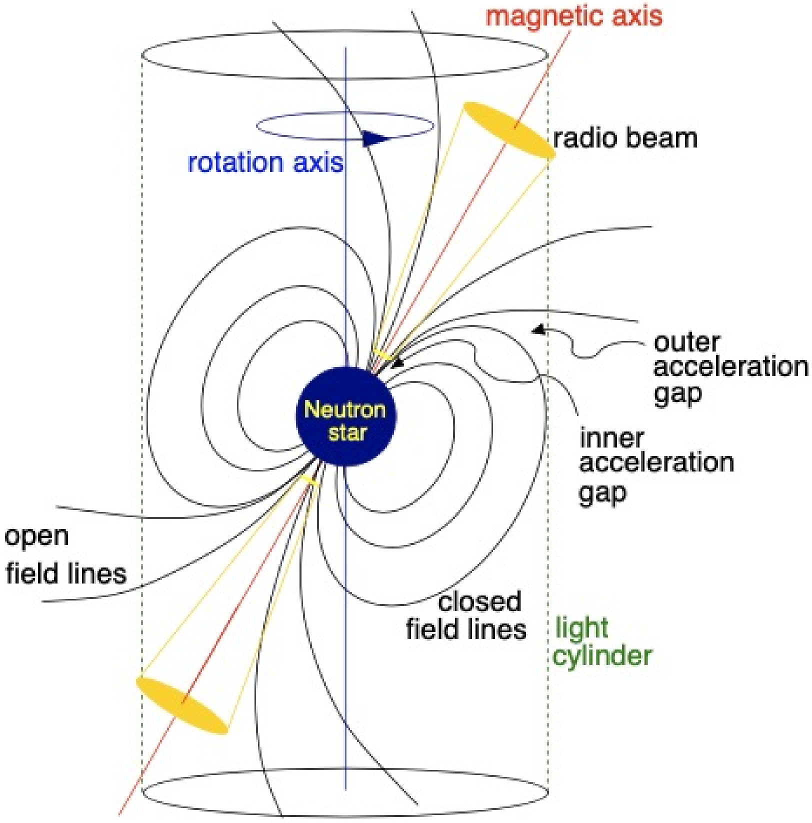

Figure 1.2: Toy model of a rotating pulsar (Lorimer and Kramer, 2004).

radiation mechanisms, including synchrotron radiation, curvature radiation, and braking radiation. Synchrotron radiation occurs as charged particles travel along helical paths within the magnetic field, while curvature radiation is emitted as particles follow the curved magnetic field lines. Braking radiation, on the other hand, takes place when particles undergo deceleration. Photons emitted through these mechanisms are guided by the magnetic field lines to form narrow radiation beams that emanate from the polar regions and propagate along the open magnetic field lines.

As the pulsar rotates, the radiation beams sweep through space much like the beam of a lighthouse. When the line of sight of Earth-based telescopes intersects with these beams, we are able to detect the signals emitted by the pulsar. The reception of these signals is synchronized with the rotation period of the pulsar. Note that because the rotational axis and the magnetic axis of the pulsar are not aligned we received periodic signals as the beams periodically cross the observer's line of sight, making the detection of the signals more pronounced. This periodic appearance of signals, akin to the sweep of a lighthouse beam across the night sky, is thus referred to as the "lighthouse effect.

1.3.2 Pulsar rotation and emission evolution

Following the introduction of the lighthouse effects of pulsars, this section addresses the rotational and emission evolution of pulsars.

The radiation emitted by pulsars affects their angular momentum, causing their rotation speed to gradually slow down. The rate of loss of rotational kinetic energy E_{rot} can be related to the period of pulsar P and its rate of increase $\dot{P}$.

$$\dot{E} \equiv -\frac{\mathrm{d}E_{\mathrm{rot}}}{\mathrm{d}t} = -\frac{\mathrm{d}\left(I\Omega^2/2\right)}{\mathrm{d}t} = -I\Omega\dot{\Omega} = 4\pi^2 I\frac{\dot{P}}{P^3} \tag{1.2}$$

where I is the moment of inertia, $\Omega = 2\pi/P$ is the angular velocity and the quantity $\dot{E}$ represents the total power output by the neutron star and can be called "spit-down luminosity".

For the pulsar rotational deceleration, we can introduce the braking index as a parameter to describe it. According to classical electrodynamics, the radiation power $\dot{E}$ of a rotating magnetic dipole radiates its energy at its rotation frequency

$$\dot{E}_{\text{dipole}} = \frac{2c^3|m|^2\Omega^4 \sin^2 \alpha}{3R^6} \tag{1.3}$$

Here α is the angle between the spin axis and magnetic moment, and $|m|$ is the moment of the magnetic dipole. Establishing the relation between Eq. (1.2) and (1.3), the evolution of the rotation frequency can represented as:

$$\dot{\Omega} = -\frac{2|m|^2 \sin^2 \alpha \Omega^2}{3Ic^3} \tag{1.4}$$

If we represent the Ω in the form of $v = 1/P$ and the Eq. (1.4) by power law

$$\dot{v} = -K\dot{v}^n \tag{1.5}$$

The break index n can be computed as a function of the constant K and the pulsar. In the ideal case, i.e. pulsar's rotational deceleration is caused entirely by magnetic dipole radiation, according to Eq (1.4) we have $n = 3$. In reality, however, there are other dissipation mechanisms, which lead us to estimate n with the the second-order derivatives of the observed v. Derivation of Eq (1.5) eliminates K and yields:

$$n = \frac{v\ddot{v}}{\dot{v}^2} \tag{1.6}$$

In this way, n is in the range of 1.4 and 2.9 by actual observation (Kaspi and Helfand, 2002; Spitkovsky and Arons, 2004).

Furthermore, starting from the Eq (1.5), we can estimate more physical quantities of the pulsar with the theory of split-down. Eq (1.5) in its expression in terms of the pulse period is:

$$\dot{P} = KP^{2-n} \tag{1.7}$$

Then the age of the pulsar can be estimated with the assumption that constant K and break index n are not equal to 1.

$$T = \frac{P}{(n-1)\dot{P}}\left[1 - \left(\frac{P_0}{P}\right)^{n-1}\right] \tag{1.8}$$

here P_0 is initial period of pulsar. If we assume that break index $n = 3$ and $P_0 \ll P$ then the estimated age of the pulsar τ_c is:

$$\tau_c = \frac{P}{2\dot{P}} \tag{1.9}$$

We can also use the spin-down theory to estimate the strength of the surface magnetic field B_S, which cannot be directly measured. Bringing relationship between magnetic moment and magnetic field strenghth ($B \approx |m|/r^3$) into Eq (1.4), we can obtain:

$$B_S = B_S(r = R) = \sqrt{\frac{3c^3}{8\pi^2}\frac{I}{R^6 \sin^2 \alpha}P\dot{P}} \tag{1.10}$$

1.3.3 Radio beam emission

In the preceding Sections 1.3.1, we have discussed that pulsars emit radiation beams from their poles. We will now delve further into the mechanism by which pulsars emit radio waves through the open magnetic field lines near their magnetic poles, resulting in the observable radio pulses.

Regarding the geometric structure and emission mechanism of pulsar radio emission, the widely accepted model posits a conical beam centered on the magnetic axis. The open field line region provides a confined space for a stable radio beam, and this model also accounts for the observed properties of pulsar profile shapes. As depicted in Figure 1.3, the angular radius ρ of the emission cone is centered on the magnetic axis and forms an angle α with the rotation axis. β is the minimum angular separation from the magnetic axis to the observer's line of sight is referred to as the impact parameter, which describes

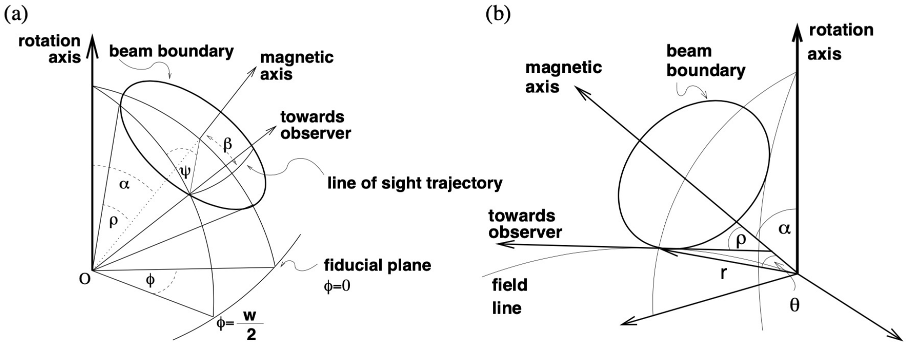

Figure 1.3: Representation the geometry of pulsar emission beams (Lorimer and Kramer, 2004).

the relative position of the observer's line of sight with respect to the pulsar's magnetic axis. The position angle of the linearly polarized emission Ψ from the pulsar (PPA) is measured with respect to the projected direction of the magnetic axis. As the pulsar rotates, the observer's line of sight traverses a curved path within the beam. The emission along this path is observed as the pulse profile, with the path length corresponding to the pulse width. The observed pulse width W, a fraction of the rotation period, can be expressed in terms of two geometric parameters relative to the cone's angular radius ρ: α corresponds to the angle between the rotation axis and the magnetic axis, and the impact parameter β represents the minimum angular distance between the magnetic axis and the line of sight (Gil et al., 1984).

$$\cos\rho = \cos\alpha\cos\alpha + \beta + \sin\alpha\sin\alpha + \beta\cos\frac{W}{2} \tag{1.11}$$

If the emission is in the region of the open magnetic field lines, the relation among the emission cone's angular radius ρ with the angle θ between the magnetic axis and the region of the emission can be expressed under the polar coordinates:

$$\tan\theta = -\frac{3}{2\tan\rho} \pm \sqrt{2 + \left(\frac{3}{2\tan\rho}\right)^2} \tag{1.12}$$

Moreover, the variation of Ψ can be elucidated by the Rotating Vector Model (RVM)(Radhakrishnan and Cooke, 1969). This model is predicated on the concept that as the pulsar rotates, the beam sweeps across the observer, causing the projected magnetic field direction to change. The observed PPA varies gradually at the outer wings of the pulse profile and rapidly at the profile's center. The RVM model predicts the swing of the PPA as follows:

$$\tan(\Psi - \Psi_0) = \frac{\sin\alpha\sin(\phi - \phi_0)}{\sin(\alpha+\beta)\cos\alpha - \cos(\alpha+\beta)\sin\alpha\cos(\phi-\phi_0)} \tag{1.13}$$

Here ϕ is the rotational phase of the pulsar, ϕ_0 is the phase at the center of the observed pulse profile, and Ψ_0 is the PPA value corresponding to ϕ_0. In theory, the RVM can be employed to estimate the geometric parameters α and β. However, in practice, the narrow width of the pulse profiles sometimes makes accurate estimation of these parameters quite challenging. Despite these difficulties, it is still possible to infer these parameters by analyzing the gradient of the PPA variation with respect to the rotational phase, particularly at the center of the pulse profile where the gradient is most pronounced (Everett and Weisberg, 2001). This approach allows for valuable inferences about the magnetic field structure and observational geometry of pulsars, even under constrained observational conditions

$$\left(\frac{d\Psi}{d\phi}\right)_{\max} = \frac{\sin\alpha}{\sin\beta} \tag{1.14}$$

which is measured with $\phi = \phi_0$.

1.3.4 Interstellar dispersion

When a pulsar emits radio signal, these signal must travel through the interstellar medium (ISM) to reach observers on Earth. The ISM contains free electrons scattered throughout space, which affect the speed of propagation of the radio waves. The propagation speed of electromagnetic waves in the medium depends not only on their frequency but also on the plasma density of the medium, that is, the density of free electrons. Specifically, electromagnetic waves of different frequencies propagate at different speeds in the presence of these electrons, a phenomenon known as dispersion effect. The refractive index u of the plasma, which measures the change in the propagation speed of electromagnetic waves in the medium relative to the speed in a vacuum, can be expressed by the following formula:

$$u^2 = 1 - \left(\frac{f_p}{f}\right)^2 \tag{1.15}$$

with f_p is the plasma frequency, defined as:

$$f_p = \sqrt{\frac{e^2 n_e}{\pi m_e}} \tag{1.16}$$

here, n_e is the electron density ($n_e \approx 0.03 cm^{-3}$), e is the elementary charge, and m_e is the electron mass, so that $f_p \approx 1.5$ kHz.

It can be concluded that as the frequency of the electromagnetic wave approaches or falls below the plasma frequency, the propagation speed of the wave slows down. In the case of high frequency, the refractive index n tends to 1, meaning the propagation speed approaches the speed of light in a vacuum. Low-frequency waves are more affected in the plasma, propagating slower than high-frequency waves. This explains the noticeable time delay of low-frequency signals from pulsars relative to high-frequency signals. Figure 1.4 presents a graphical representation of the dispersion effect observed on pulsar `PSR B1356-60`.

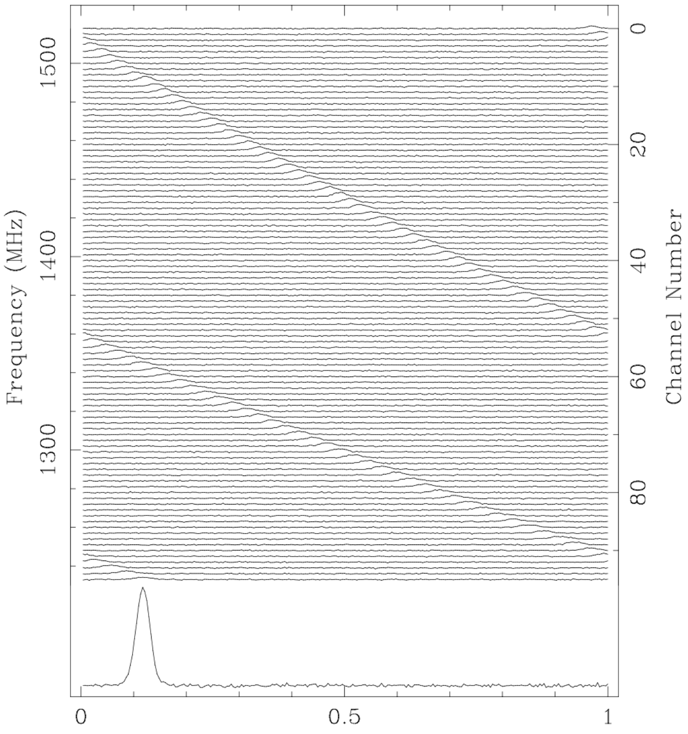

Figure 1.4: Illustration the dispersion effect on pulsar signals as they propagate through the interstellar medium, while Y-axis is frequency of observation, the X-axis spans one period pulsar. The observation of pulsar `B1356-60` with dispersion measure at 295 cm^{-3}pc (Lorimer and Kramer, 2004).

This phenomenon is of great importance in pulsar astronomy. By measuring the time delay of signals at different frequencies (referred to as dispersion delay), we can estimate the total number of free electrons between the pulsar and Earth (i.e., the dispersion measure DM), thereby inferring the distance to the pulsar and the structure of the interstellar medium. The time delay between two frequencies f_1 and f_2 can be calculated by:

$$\Delta t(f_1, f_2) = \mathcal{D} \cdot \left(\frac{1}{f_1^2} - \frac{1}{f_2^2} \right) \cdot \int_0^D n_e(l) dl \tag{1.17}$$

where $\mathcal{D}$ is the dispersion constant defined as:

$$\mathcal{D} = \frac{e^2}{2\epsilon_0 m_e c} \approx 4.148808 \pm 0.000003 \times 10^3 \mathrm{MHz}^2 \mathrm{pc}^{-1} \mathrm{cm}^3 \mathrm{s} \tag{1.18}$$

The dispersion measure is defined as the integral of the column density of free electrons along the line of sight. This total number of free electrons along the path through which the electromagnetic wave propagates is defined as :

$$DM = \int_0^d n_e dl \tag{1.19}$$

here n_e is the number density of free electrons in the interstellar medium, with units of cm^{-3}, d is the distance from the pulsar to the Earth, and dl is the small distance element along the path with units of parsecs (pc).

1.3.5 Population of pulsars

Since the discovery of the first pulsar in 1967, more than three thousand pulsars have been observed, all recorded in the PSRCAT catalogue[3]. These pulsars exhibit a wide range of rotation periods, with the fastest being around 0.001 seconds and the slowest reaching tens of seconds. As shown before in the Section 1.3, the rotation period of a pulsar is not constant and it slows down gradually over time. This deceleration can be measured by observing changes in the period of the pulsar and is shown in the Figure 1.5, which shows the relationship between the period of the pulsar and its period derivative.

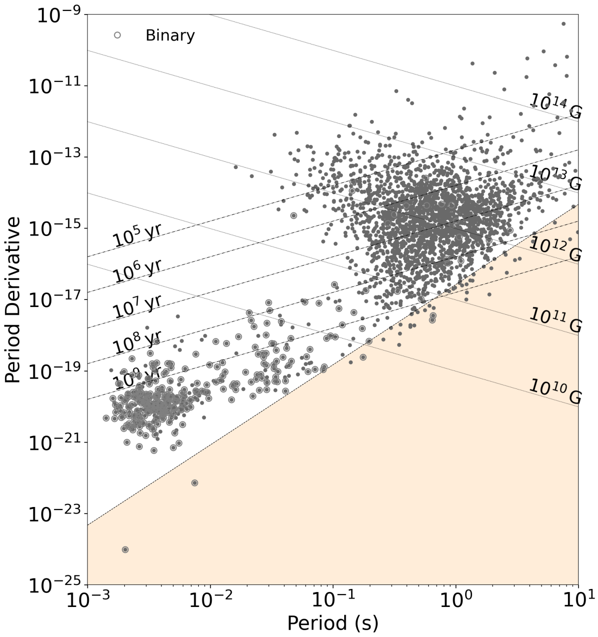

Figure 1.5: The diagram $P - \dot{P}$ of pulsar, illustrating the relationship between its period and the variation of its period. The figure is generated by the code of Pitkin (2018).

---

[3] https://www.atnf.csiro.au/research/pulsar/psrcat/

The evolution of a pulsar typically begins with a supernova explosion, resulting in the formation of a young pulsar. These young pulsars have very fast rotation rates, with periods typically ranging from tens of milliseconds to a few seconds. For example, the famous Crab Pulsar (`PSR B1509-58`) is a young pulsar with a rotation period of about 0.033 seconds. Over time, their rotation slows down, shown in the $P - \dot{P}$ diagram as a movement towards the centre and lower right. Some pulsars may eventually cross the so-called 'death line' and enter the yellowish region in the lower right-hand corner of the diagram, meaning that their rotation has slowed down to the point where they can no longer sustain effective radio emission, and thus become invisible in the radio bands.

However, if the pulsar is member of a binary system, it accretes material from its companion star and undergo a process known as recycling to become a millisecond pulsar. During this process, the accreted material not only increases the angular momentum of the pulsar, thereby causing a faster rotation, but may also interact with the magnetic field to weaken it, which in turn reduces the braking torque. Finally, if conditions are suitable, the rotation of the pulsar can reach the millisecond level, characterising it as a millisecond pulsar. For example, `PSR J0437-4715` is a typical millisecond pulsar with a rotation period of only 5.76 milliseconds. Millisecond pulsars, with rotation periods usually between 1-10 milliseconds, form a distinct group in the lower left corner of the $P - \dot{P}$ diagram, clearly set apart from normal pulsars. Figure 1.6 provides a schematic of the evolution of a pulsar companion.

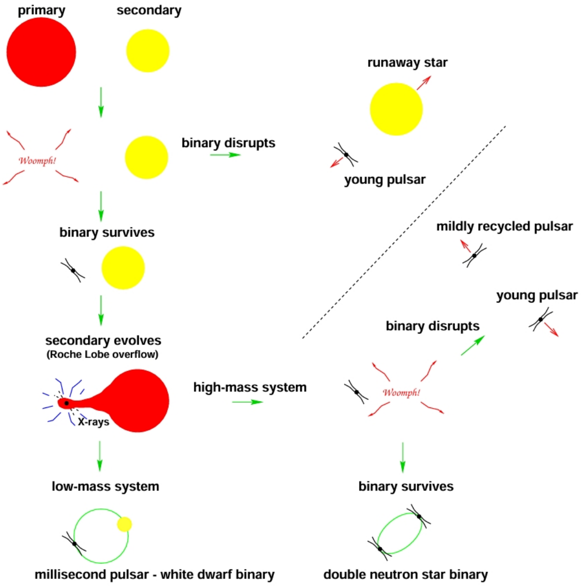

Figure 1.6: Schematic representation of the evolution of pulsar binaries (Lorimer, 2008).

This process can take place over a timescale of tens of millions to billions of years and represents a fundamental phase in the evolutionary history of pulsars. The study of these millisecond pulsars enables scientists to gain a deeper understanding of the physical properties of pulsars and their evolutionary history in the universe. Notable research in this field includes detailed observations of the Crab Pulsar (Weisberg and Taylor, 2004), as well as the discovery of the millisecond pulsar `PSR J1748-2446ad` (Hessels et al., 2006), which has a period of 1.4 milliseconds and is currently the fastest-spinning pulsar detected.

Résumé du chapitre

Les pulsars sont des restes d'explosions de supernovae qui sont des étoiles à neutrons hautement magnétisées et en rotation rapide. Depuis leur découverte en 1967 (Hewish et al., 1968), ils jouent un rôle central dans la recherche astronomique. Ils servent de laboratoires naturels pour tester la physique fondamentale et servent également de détecteurs d'ondes gravitationnelles. Grâce à leurs périodes de rotation incroyablement stables, les pulsars sont réputés pour être les horloges les plus précises de l'univers, ce qui est significatif pour une variété d'études astrophysiques.

La quête pour comprendre les pulsars a commencé au début du 20ème siècle avec les efforts pour comprendre la composition du noyau atomique. Cette recherche a conduit à la découverte du neutron par Chadwick en 1932 (Chadwick, 1932), préparant le terrain pour le concept d'étoile à neutrons. Les avancées théoriques de Baade et Zwicky en 1934 (Baade and Zwicky, 1934), ainsi que les travaux ultérieurs de Landau, d'Oppenheimer et Volkoff (Oppenheimer and Volkoff, 1939), ont jeté les bases pour comprendre ces corps célestes denses, aboutissant à la prédiction de la limite de Tolman-Oppenheimer-Volkov (TOV) pour les masses des étoiles à neutrons.

L'effet de phare des pulsars, où leur rotation rapide balaye leurs faisceaux de rayonnement à travers le ciel, illustre la manière dont ces étoiles à neutrons émettent des radiations. Le ralentissement graduel de leur rotation dû à l'émission de rayonnement peut être modélisé mathématiquement pour estimer leur âge et l'intensité de leur champ magnétique superficiel.

Le mécanisme d'émission des impulsions radio des pulsars est également étudié. Le modèle accepté suggère que les ondes radio sont émises à travers des lignes magnétiques ouvertes près des pôles magnétiques, qui forment des impulsions radio observables. La structure géométrique et le mécanisme d'émission de ces impulsions, y compris l'angle de vue conique et la ligne de visée de l'observateur, sont cruciaux pour comprendre les propriétés observées des profils des pulsars.

La dispersion interstellaire, qui correspond à la variation de la vitesse de propagation des ondes électromagnétiques dans le milieu interstellaire, est essentielle à l'étude des pulsars. En mesurant le délai de temps des signaux à différentes fréquences, la mesure de dispersion peut être estimée, fournissant des informations sur la distance des pulsars et la structure du milieu interstellaire.

Un aperçu de la population de pulsars révèle une large gamme de périodes de rotation, allant de 0,001 seconde pour les plus rapides à plusieurs secondes pour les plus lents. L'évolution des pulsars commence généralement par une explosion de supernova, qui conduit à la formation de jeunes pulsars avec des vitesses de rotation très élevées. Avec le temps, leur rotation ralentit et certains deviennent invisibles dans les bandes radio. Cependant, ceux qui appartiennent à des systèmes binaires peuvent devenir des pulsars millisecondes grâce à un processus appelé recyclage, qui augmente leur moment angulaire et accroît leur vitesse de rotation.

Chapter 2

Radio astronomy telescope – NenuFAR

2.1 Radio astronomy and telescopic techniques

Radio astronomy is a scientific discipline that employs the use of radio waves for the observation and study of the universe. Unlike traditional optical astronomy, which relies on electromagnetic waves in the visible band, radio astronomy fields celestial objects and phenomena through the detection of radio frequency emissions. The development of radio telescopes can be traced back to the early 20th century when scientists began to explore the application of radio waves in astronomy. In 1933, American engineer Jansky (1933) first detected radio signals from the center of the Milky Way, marking the birth of radio astronomy. Subsequently, radio telescope technology has evolved from early single-antenna designs to complex interferometric arrays. During the 1960s, radio telescopes began to employ aperture synthesis techniques, significantly enhancing the resolution of observations (Ryle and Hewish, 1960). By the 21st century, the design and construction of radio telescopes had reached an unprecedented scale and complexity, with notable examples including the Atacama Large Millimeter/sub-millimeter Array (ALMA) operating at high frequencies (Wootten and Thompson, 2009; Mignani et al., 2017), the Low frequency array telescope (LOFAR) (van Haarlem et al., 2013), and the upcoming Square Kilometre Array (SKA) (Dewdney et al., 2009), which will be a vast array spanning a wide range of frequencies.

The observational targets of radio astronomy include, but are not limited to, stars, galaxies, black holes, pulsars, the interstellar medium, and the cosmic microwave background radiation. The signals emitted by these objects and phenomena in the radio wavelength band provide invaluable insight into the structure, composition, and evolution of the universe. Since radio wavelengths are longer than those of visible light, they are not absorbed by the atmosphere and can penetrate interstellar dust and gas to reach the surface on Earth.

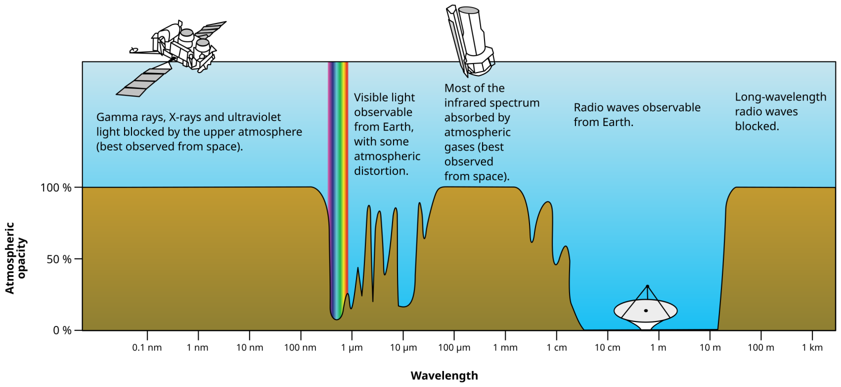

Figure 2.1: Spectral window of radio waves detectable from Earth, highlighting Earth's atmospheric absorption and scattering of different electromagnetic wavelengths (picture extracted from Wikipedia[1]).

---

[1] By NASA (original) File:Atmospheric electromagnetic transmittance or opacity.jpg., Public Domain, https://commons.wikimedia.org/w/index.php?curid=5577513

As shown in Figure 2.1, the frequency range involved in radio astronomy typically spans from 30 MHz to 300 GHz, covering bands from Very High Frequency (VHF, 30 MHz - 300 MHz, corresponding to wavelengths of 10 m to 1 m) to microwave (3 GHz-300 GHz, corresponding to wavelengths of 10 cm to 1 mm). This broad frequency range allows radio astronomers to study a wide variety of astrophysical processes. For instance, lower frequency radio waves (typically in the VHF to UHF range, 30 MHz to a few GHz) are commonly used to study pulsars and supernova remnants, while higher frequencies (such as microwaves, 3 GHz-300 GHz) are used to observe molecular clouds, star-forming regions, and active galactic nuclei. By analyzing radio signals at different frequencies, scientists can obtain information about the temperature, density, chemical composition, and magnetic fields of celestial bodies.

Radio telescopes typically consist of one or more large antennas designed to capture the faint radio signals from the cosmos. The design and size of these antennas depend on the observed band and the target being studied. A radio telescope can be a single large-diameter antenna or an array of several smaller antennas. Arrays of multiple antennas use interferometry to simulate a virtual large-diameter antenna with higher resolution. The received radio wave signals are first converted into electrical signals, which are then amplified, filtered, and digitized. These data are then processed and analyzed by a computer to produce images or spectral data of celestial objects.

2.2 NenuFAR

NenuFAR (New Extension in Nançay Upgrading LOFAR) is a low-frequency radio telescope located at the Nançay Radio Observatory in France, specifically designed for radio astronomical observations in the low-frequency range ($10-85$ MHz). As part of the LOFAR network, NenuFAR can function independently for scientific observations, while also integrating with other LOFAR stations to offer higher sensitivity and spatial resolution. The construction of NenuFAR began in 2014, and it officially started scientific operations in 2019. This project was developed in collaboration with multiple institutions, including the French National Centre for Scientific Research (CNRS) and the Paris Observatory (Observatoire de Paris), with the aim of advancing research in low-frequency radio astronomy.

What sets NenuFAR apart is its focus on a wide range of applications in the low-frequency radio spectrum, from precise measurements of pulsars to the detection of signals from the cosmic dawn. Its design aims to fill the observational gaps in the low-frequency coverage of LOFAR, thereby improving our understanding of the radio universe.

NenuFAR, similar to LOFAR and the upcoming SKA, is a radio astronomical array composed of many small antenna elements. These elements function in conjunction through interferometric techniques to facilitate high-resolution observations, the detection of extremely faint signals, and the execution of rapid surveys. In the design and operation of such arrays, the configuration of antennas and the process of calibration represent pivotal stages that directly influence the quality of the observational data.

Recent research, such as studies utilizing regularized maximum likelihood estimation frameworks, provides new insights for radio interferometric imaging (Mhiri et al., 2024). Other works suggest the barankin-type bound (BTB) (Wang et al.) as a more appropriate array design criterion for low SNR. Furthermore, algorithms like the parallel multi-wavelength calibration algorithm (PMCA) (Brossard et al., 2018) and the space alternating generalized expectation-maximization (SAGE) (Mhiri et al., 2022) approach for calibrating radio interferometric arrays in the presence of RFI, contribute significantly to the advancement of the field.

2.2.1 NenuFAR's mini-array and receiver

NenuFAR is a phase-array radio telescope, with its core comprising a compact antenna array made up of 96 mini-arrays (MA, hereafter), each consisting of 19 individual antenna elements. This configuration results in a total of 1824 antenna elements across a circular area with a radius of approximately 400 meters. In addition to the 96 MA that constitute the core of NenuFAR, an additional 6 MA are projected to be installed approximately 3 kilometers from the central array. This expansion is intended to enhance the telescope's imaging capabilities. Figure 2.2 presetnts the distribution of NenuFAR's mini-arrays.

The MA of NenuFAR consist of 19 crossed dipoles arranged in a hexagonal pattern, each elevated 1.6 meters above a 3-meter by 3-meter metallic ground plane. Oriented at a 45-degree angle to the meridian, these dipoles are categorized into two linear polarizations: northeast-southwest (NE or X dipoles) and

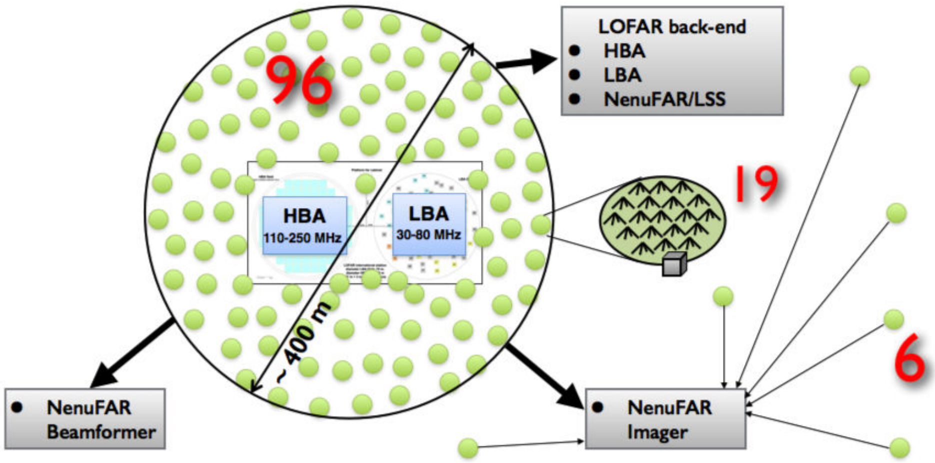

Figure 2.2: Schematic distribution diagram of NenuFAR's mini-arrays. The distribution is Gaussian in the radial direction and uniform in the angular direction (picture ectracted from NenuFAR webpage[2]).

northwest-southeast (NW or Y dipoles). The overall diameter of a single Mini-Array spans approximately 25 meters. These Mini-Arrays synthesize signals through an analog phase and summation system, which creates a steerable beam that enhances sensitivity within a specific field of view. The analog phasing system utilizes variable cable lengths for delay lines, enabling precise pointing towards 16,384 distinct directions across the sky, from the horizon to the zenith. After the analog processing, each Mini-Array transmits two linearly polarized signals (NE and NW) to the receivers via coaxial cables, allowing for the capture of the full polarization state of incoming electromagnetic waves. This configuration ensures high sensitivity of NenuFAR over the 10-85 MHz frequency range, which is mainly limited by sky noise.

NenuFAR's receiving system comprises the main standalone receiver, LaNewBa (LOFAR super station Advanced NEW Backend), which digitizes the analog signals from the 96 MA across two polarizations at a sampling frequency of 200 MHz. The signals are then organized into 1024 sub-bands, each with a bandwidth of 195.3125 kHz, and processed through beamforming to create 768 beamlets. These beamlets, defined by a triplet (f_c, θ, ϕ) of central frequency, sky direction, and polarization, allow for highly sensitive and directed observations across an instantaneous bandwidth of 150 MHz. These 768 beamlets can be distributed in (f_c, θ, ϕ) as desired such as 2 beams with 75 MHz bandwidth (full-bandwidth) or 4 beams with 37.5 MHz (half-bandwidth). The beam data produced by LaNewBa will be sent to the NenuFAR back-end processing system. The data processing flowchart of the NenuFAR receiver is depicted in Figure 2.3.

2.3 NenuFAR in pulsar observation

For pulsar observations (Bondonneau et al., 2021), the data are beamformed by LaNewBa and then delivered via WLAN to the calculator UnDySpuTed (Dynamic spectroscopy & pulsars modes) for next step processing. UnDySPuTeD, comprised of two identical computers, operates *(a)* the LUPPI (Low frequency Ultimate Pulsar Processing Instrumentation) software specifically for real-time processing of pulsar observation data, including high data rate handling and real-time coherent dedispersion across the full frequency range of NenuFAR and *(b)* the **tf** (time-frequency) software to build dynamic spectra with different resolutions. UnDySPuTeD can provide support for pulsar observations in four data modes: (i) *Waveform mode*, charged by **tf**; (ii) *Folded mode*, provided by LUPPI; (iii) *Single-pulse mode*, also provided by LUPPI; (iv) *Dynamic spectrum mode*, charged by **tf**.

Waveform mode – In waveform mode, NenuFAR records the raw complex voltage data (195312 pairs of complex X and Y signals values per second per beamlet) directly to disk without processing by the computational threads. This mode retains the most complete observation data, including full time, frequency,

---

[2] NenuFAR webpage: https://nenufar.obs-nancay.fr/en/astronomer/

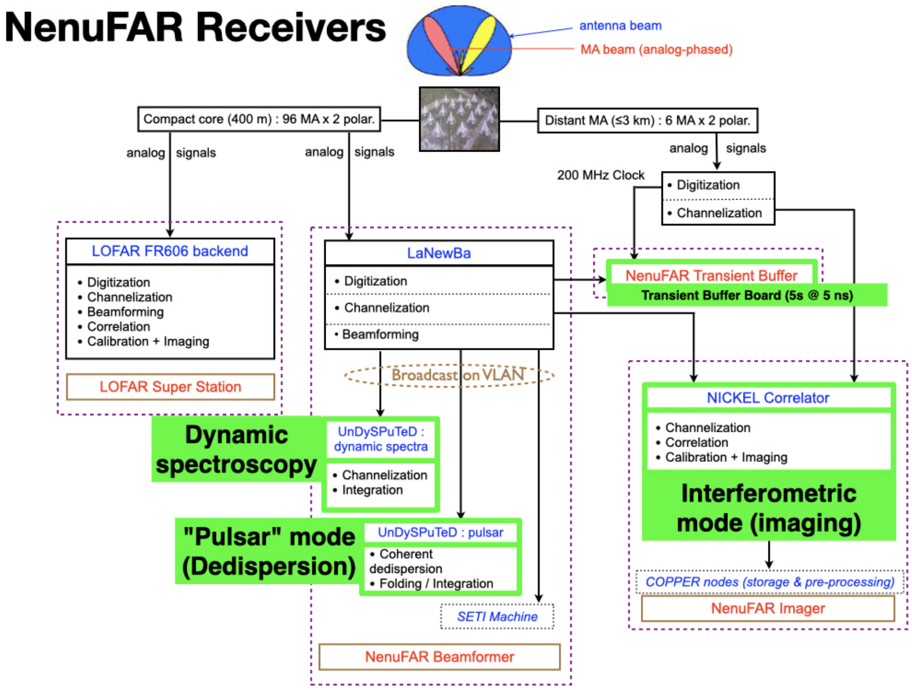

Figure 2.3: Schematic diagram of NenuFAR receiver (picture ectracted from NenuFAR webpage).

and polarization information, but correspondingly generates very large amounts of data. This model is useful for studying the raw properties of pulsar signals,but requires significant data storage and further data processing power.

Folded mode – The folded mode creates the periodic profile of a pulsar by time-folding the signal after de-dispersion. This mode is commonly used for routine observations of pulsars and can provide detailed information about the radiation characteristics of pulsars, such as the shape of the pulse profile, polarization characteristics, and changes in frequency over time.

Single-pulse mode – Single-pulse mode allows researchers to analyze individual pulse events rather than folding multiple pulses. This mode is very useful for studying inter-pulse variations and transient phenomena in pulsars. In this mode, the accuracy of dispersion measurement (DM) is crucial for distinguishing signals from different pulses. By precisely measuring the dispersion delay, researchers can better understand the emission mechanisms of pulsars and the propagation characteristics of the signals.

Dynamic spectrum mode – The dynamic spectrum mode employs GPU acceleration to process the raw observational data, providing a detailed two-dimensional representation of the pulsar signal's evolution with respect to both time and frequency. This capability is crucial for examining the rapid temporal and spectral variations in pulsar emissions. The mode is specifically designed to capture transient phenomena within pulsar signals, including scintillations and profile modulations. Offering superior temporal resolution, the dynamic spectrum mode is not only valuable for analyzing the dynamic characteristics of pulsars but also optimal for executing blind surveys aimed at discovering previously unknown pulsars, in addition to study other sources such as planets, fast radio bursts, etc... This mode is also the focus of the work reported in this manuscript. More precisely a framework to simulate data according to this mode will be presented in subsequent Chapter 4 and Chapter 7 describes the performance of the proposed approach when applied to dynamic spectra.

Since the start of early science operations in July 2019, NenuFAR has detected a total of 183 pulsars, 100 of which were first detected below 100 MHZ. Among these detections, 13 were millisecond pulsars, including 7 that were first detected below 100 MHz.

Résumé du chapitre

NenuFAR (Zarka et al., 2018), extension de l'interféromètre LOFAR, est un télescope radio à basse fréquence spécialisé situé à Observatoire radioastronomique de Nançay en France. Conçu pour explorer le spectre de fréquence de 10 à 85 MHz, NenuFAR complète le réseau LOFAR, en fonctionnant de manière autonome ou en synergie avec d'autres stations LOFAR pour augmenter la sensibilité et la résolution spatiale. Depuis le début de ses opérations scientifiques en 2019, après le début de sa construction en 2014, NenuFAR a considérablement contribué au domaine de l'astronomie radio à basse fréquence, comblant les lacunes d'observation de LOFAR dans cette bande de fréquence et élargissant notre compréhension de l'univers radio.

L'architecture innovante du télescope comprend un ensemble compact de 96 mini-ensembles, chacun composé de 19 dipôles croisés, et six mini-ensembles supplémentaires prévus pour être installés à 3 kilomètres du centre, afin d'améliorer les capacités d'imagerie. Ces mini-ensembles, orientés à un angle de 45 degrés et élevés au-dessus d'un sol en métal, sont sensibles à deux polarisations linéaires et, grâce à un système de phase et de sommation analogique, peuvent se concentrer sur 16 384 directions distinctes dans le ciel. Ce système, associé au récepteur principal LaNewBa, numérise les signaux à une fréquence d'échantillonnage de 200 MHz, les répartit en 1024 sous-bandes, chacune ayant une bande passante de 195,3125 kHz, et forme 768 faisceaux pour des observations hautement sensibles et dirigées.

Dans le contexte de l'observation des pulsars (Bondonneau et al., 2021), NenuFAR utilise le système UnDySpuTed, qui traite les données en quatre modes : forme d'onde, impulsion empilée, impulsion simple et spectre dynamique. Le mode en forme d'onde enregistre les données brutes, le mode impulsions empilée crée des profils périodiques, le mode impulsion simple analyse des événements d'impulsion individuels, et le mode spectre dynamique, qui utilise l'accélération GPU, offre une représentation bidimensionnelle détaillée de l'évolution des signaux des pulsars en fonction du temps et de la fréquence. Ce mode est crucial pour examiner les variations rapides des émissions des pulsars et pour mener des enquêtes visant à découvrir de nouveaux pulsars.

Les contributions de NenuFAR à l'astronomie des pulsars ont été remarquables depuis le début de ses opérations scientifiques, avec la détection de 183 pulsars, dont 100 ont été détectés pour la première fois en dessous de 100 MHz et 13 pulsars millisecondes, dont sept ont été détectés pour la première fois dans la même plage de fréquence. Ces réalisations soulignent le rôle de NenuFAR dans l'avancement de nos connaissances sur les pulsars et son potentiel dans l'étude des ondes gravitationnelles et d'autres phénomènes astrophysiques.

Chapter 3

Radio Frequency Interference (RFI)

3.1 RFI

RFI refers to the unwanted signals that contaminate the radio frequency spectrum, which is used for various astronomical observation (Baan, 2011; An et al., 2017). In the context of radio astronomy, RFI is usually caused by a variety of radio signals generated by human activities, including terrestrial communication networks such as FM radio, television broadcasting and mobile phone services, which typically operate in specific frequency bands and can generate narrow-band or broad-band interference. Satellite communications systems, such as the Global Positioning System (GPS), and aircraft communications and navigation equipment also transmit signals on their specific frequencies. In addition, terrestrial activities around the observatory, such as lighting systems, power lines and electronic fences, may produce transient or persistent infections. Figure 3.1 shows several examples of typical RFI waveforms in the time domain.

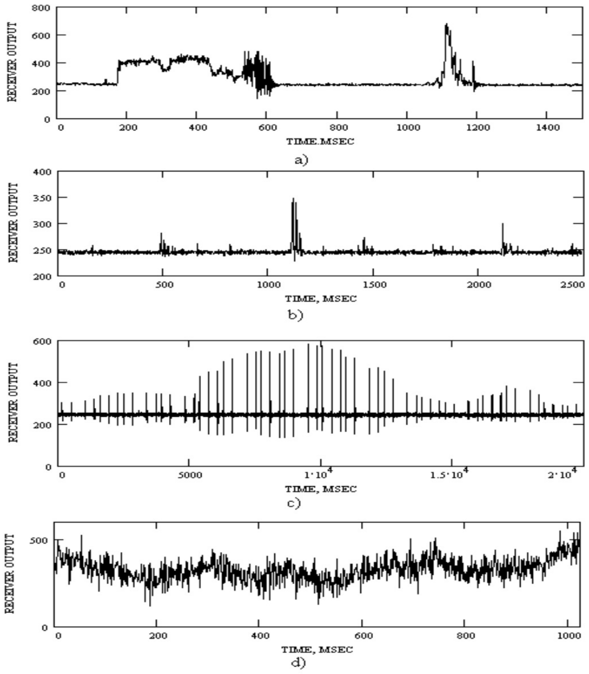

Figure 3.1: Examples of RFI waveform temporal: panels a) and b) demonstrate pulse-like RFIs, which are characterized by brief but intense signal spikes in the time domain. Panel c) Shows a radar pulse, which is an RFI that usually has a specific repetitive pattern and can be periodic or follow a regular pattern. Panel d) Demonstrates a narrow-band RFI, which occupies a narrow frequency range in the frequency domain, but may manifest itself as a persistent interference signal in the time domain (Fridman and Baan, 2001).

3.2 Importance and challenges of RFI mitigation

The presence of RFI poses a huge challenge to radio astronomy as it can disturb radio telescope in capturing weak astronomical signals (RA, 2003). Such interference may not only lead to misinterpretation of astronomical phenomena, but also rise to erroneous detection results. Radio telescopes, as a highly sensitive observational tool, are particularly vulnerable to RFI, which can mask signals from the distant universe, leading to data loss and degradation in the quality of the observations and compromising the integrity of the data collected. Therefore, the management and mitigation of radio frequency interference is a key issue in ensuring the reliability of astronomical observations and requires effective strategies to address it.

In the field of radio astronomy, the current landscape of RFI mitigation presents a multifaceted set of challenges that demand innovative solutions and collective efforts. Two principal challenges confront radio astronomy. One arises from the growing prevalence of wireless communication technologies, notably the advent of 5G and the heightened utilisation of satellite systems, which have introduced intricate and pervasive forms of RFI. These emerging sources of interference have a broader impact and exhibit complex traits that are increasingly difficult to predict and eliminate. The other aspect can be attributed to advancements in detection technology, which have led to the development of new generations of radio telescopes with enhanced sensitivity, enabling the detection of astronomical signals that were previously undetectable (Ellingson, 2005). This development provides astronomers with the ability to observe new types of objects or phenomena that were previously undetected. Furthermore, for known objects and phenomena, the capacity to detect fainter signals can facilitate more precise measurements. This is particularly evident in the detection of fast radio bursts (FRBs) and pulsars.

3.3 RFI mitigation at various observation stages

Since RFI can lead to data loss and quality degradation of astronomical observations, RFI mitigation has been taken into consideration at all stages of the observation to cope with this problem. A comprehensive overview of the methods and techniques used to mitigate RFI, with a categorization of the methods according to the stage of observation, was provided by (Baan et al., 2004a; An et al., 2017). In addition these two papers also summarized various types of RFI and their characteristics. In reviewing the existing methods, Fridman and Baan (2001) emphasise that there is no one universal RFI mitigation method and states that the best mitigation strategy should be tailored to the specific radio telescope, the type of observation, the stage of observation applied and the characteristics of the RFI environment. Based on the processing of data by the telescope, the primary approaches for RFI management can be generally categorized into 4 different stages: *i)* observatory prevention, *ii)* pre-detection strategies, *iii)* pre-correlation techniques, *iv)* post-correlation methods.

3.3.1 Observatory prevention

The first and most powerful method is to control incoming RFI signals before they enter the receiver (Baan, 2011). Prevention of RFI before the observation can can be categorized into two aspects, one being spectrum management, i.e., reserving the necessary frequency bands for astronomical observations, and restricting or prohibiting potentially interfering activities only in those bands. The second aspect is the selection of the location of the observatory, which reduces interference by locating the observatory away from man-made sources of infection.

Spectrum management – Effective spectrum management is pivotal in the ongoing battle against Radio Frequency Interference (RFI) in radio astronomy. The International Telecommunication Union (ITU) provides recommendations and guidelines that support the establishment of RFI-free zones and the implementation of regulatory measures to govern frequency usage around observatories. For instance, the ITU's Radio Regulations allocate specific frequency bands for radio astronomy observations. According to the ITU-R Report RA.2099 (RA, 2003), the preferred frequency bands for high-precision radio pulsar timing observations, which are crucial for accurate timekeeping, include the Radio Astronomy Service (RAS) range of 1400 MHz to 1427 MHz, the 406.1 MHz to 410 MHz band, the 608 MHz to 614 MHz band, and in certain instances, the 2690 MHz to 2700 MHz band. These allocations are part of the ITU's broader spectrum management strategies designed to minimize interference and maximize the sensitivity of radio astronomical observations. Moreover, spectrum management strategies include proactive identification and coordination with active frequency users to prevent RFI incidents, as well as continuous monitoring of the radio spectrum to detect and address new sources of interference.

Telescope site selection – The most effective initial strategy for RFI mitigation is the selection of an observatory site that is naturally free from RFI. This necessitates the identification of locations exhibiting minimal human activity and radio frequency usage. The choice of an optimal site for a radio telescope is a crucial decision that necessitates the balancing of a number of factors that are essential for successful astronomical observations. An optimal location can provide an unpolluted radio environment, free of human-made noise and radio frequency interference, which is common in remote areas. Furthermore, it can facilitate the establishment of radio quite zone around the telescope and control the use of specific frequency bands in the zones. For example, NenuFAR telescope introduced in the Section 2.2, located approximately 200 kilometers south of Paris, is nestled in the heart of the Sologne forest at the Nançay Observatory, an area with minimal human activity and designated as a radio quiet zone to protect against electromagnetic interference. This protection zone is illustrated in Figure 3.2. While the next generation of astronomical telescopes, the SKA, is being built in Australia and South Africa, as they have the least radio interference.

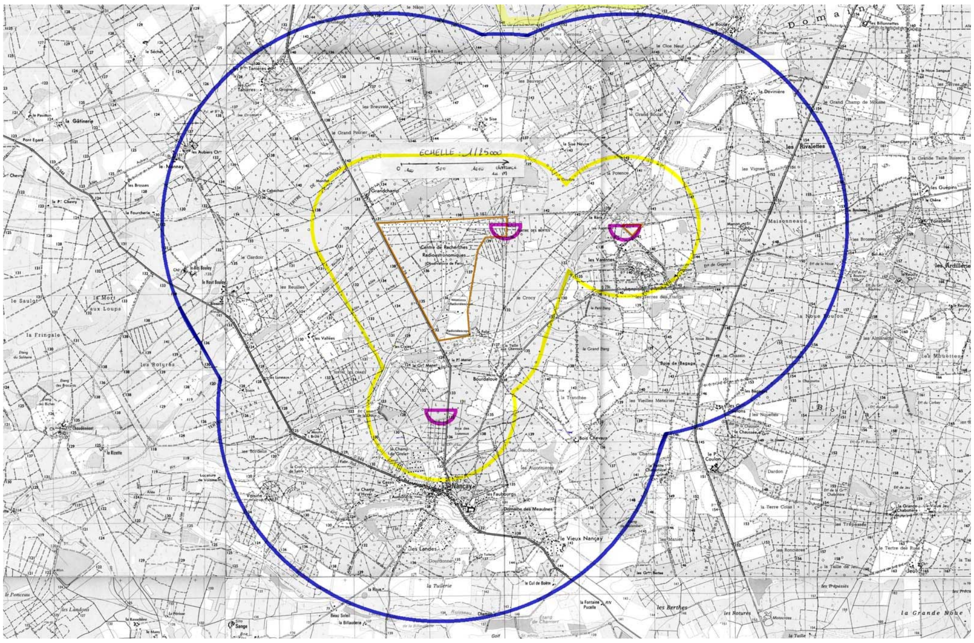

Figure 3.2: Map of protection zones for the Nançay radio observatory, radio quiet protection zone (in blue and yellow) and extent and nature of planned servitudes against obstacles (in magenta) (pciture extracted from website of Nançay radio observatory[1]).

3.3.2 Pre-detection strategies

In radio astronomy, pre-detection strategies are essential to minimise radio frequency interference before it affects the data acquisition process. These strategies are implemented at the receiver level and are designed to filter out known strong RFI signals outside the telescope operating band.

Bandpass and high/low pass filters – A standard pre-detection method involves the installation of bandpass or high/low pass filters within the receiver system. These filters help to excise RFI in the frequency domain, preventing strong signals from outside the desired frequency bands from entering the data stream. While this method is effective, it may result in an insertion loss and raise the system temperature at frequencies close to the band-edge, potentially leading to data loss for continuum observations.

---

[1] Website of Nançay radio observatory: https://www.obs-nancay.fr/protection-radio/

Blanking or data acquisition process stopping – Another pre-detection strategy is to halt the data-taking process, such as stopping the accumulation of data in the correlator, during periods of strong RFI reception. This method is particularly useful for dealing with impulsive and periodic RFI signals and helps to avoid data loss by only stopping the process during specific time windows.

3.3.3 Pre-correlation techniques

The pre-correlation stage is identified as a critical point in the data acquisition process where RFI mitigation techniques can be effectively applied before the data undergoes correlation, which is the processes of combining signals from multiple antennas to enhance the sensitivity and resolution of the observation. Common techniques applied in this stages to reduce RFI including real-time digital processing, adaptive noise cancellation (ANC), and spatial filtering.

Real-time digital processing – For single-dish radio telescopes and array instruments, real-time digital processing can be implemented as part of the Intermediate Frequency (IF) processing. This method is effective for impulsive RFI and requires fast data sampling and the availability of computing resources. Real-time digital processing technologies, such as filtering based on FFT, can distinguish between astronomical signals and RFI, allowing for the effective removal of interference before it affects data quality. The system RFIMS proposed by Baan et al. (2004b) utilizes an experimental Field-Programmable Gate Array (FPGA) configuration for real-time digital signal processing,which initially carries out frequency domain filtering. Subsequently, it applies a thresholding process to the incoming data, employing a technique known as cumulative summing to manage RFI effectively. Moreover, by employing algorithms such as adaptive filtering and threshold processing (Sclocco et al., 2019), observatories can efficiently handle time-variable and persistent RFI, minimizing data loss and targeting only the RFI-contaminated parts of the data for processing while preserving the surrounding clean data for further analysis.

Adaptive noise cancellation – Adaptive noise cancellation leverages the principle of adaptive filtering, which involves transforming the incoming data into the frequency domain using a Fast Fourier Transform (FFT), applying an adaptive algorithm to identify and attenuate the interference, and then reverting the data back to the time domain through an inverse FFT. The effectiveness of ANC is particularly pronounced when dealing with RFI that has a significant interference-to-noise ratio (INR), where the interference dominates the system noise. This method is based on Wiener filtering (Barnbaum and Bradley, 1998) and can achieve suppression of the interfering signal to a level approximately equal to its instantaneous INR. ANC is well-suited for applications where spectral information is not critical, such as in pulsar studies (Kesteven et al., 2005) and continuum observations. In practice, ANC can be implemented through different strategies including subtracting a reference data-channel from the signal data-channel using a copy of the RFI. This method compares signals with and without RFI and may utilize an auxiliary antenna pointed at the interference source. Additionally, parametric estimation and subtraction of known RFI signals have successfully removed specific satellite transmissions by leveraging their modulation properties, achieving substantial signal cancellation without extra antennas (Ellingson et al., 2000; Briggs et al., 2000).

3.3.4 Post-correlation methods

The post-correlation stage involves advanced digital signal processing techniques, including flagging and excision, to identify and remove RFI from the data. This is a more intensive process compared to pre-correlation stage, as the data are often integrated over various time intervals, which can range from the sampling time up to several seconds. The RFI mitigation can be implemented through different methods in this stage, such as digital filtering of signals outside the target frequency range, spatial nulling, adaptive cancellation of RFI using reference channels or antennas, and the application of thresholding algorithms. Post-correlation RFI mitigation employs a suite of methods that, while also relevant to the pre-correlation stage, serve different purposes at each stage. During pre-correlation, the focus is on preemption—guarding the data stream against the initial incursion of RFI to prevent contamination from the outset. This early intervention is key to maintaining the purity of the dataset and reducing the need for extensive downstream mitigation. Conversely, the post-correlation stage is tasked with the detection and rectification of RFI that has either gone undetected by earlier measures or has become more discernible following the integration of data. This phase is inherently more analytical and demanding, requiring a thorough review of the integrated data to identify and address RFI occurrences that surface at this later stage. The post-correlation efforts are more likely to result in data loss due to the broader impact of RFI on the consolidated dataset. Moreover, this stage demands a heavier reliance on computational power to execute complex algorithms and possibly leverage machine learning technologies. These advanced tools are crucial for effectively dissecting the integrated data and achieving a higher success rate in RFI removal. In essence, post-correlation RFI mitigation is a deeper, more resource-intensive process that builds upon the foundational work of pre-correlation strategies. It represents the next line of defense in the ongoing effort to ensure the accuracy and reliability of radio astronomical data. This manuscript focuses on this stage to investigate and develop a new algorithm for removing RFI from the observations of a new generation telescope like NenuFAR, so the Section 3.4 will review the various algorithms that have been applied in this stage, including traditional algorithms based on signal processing or thresholding as well as algorithms based on advanced machine learning or deep learning.

3.4 Overview of post-correlation RFI mitigation methods

A variety of methods have been put forth with the aim of mitigating the RFI during the post-correlation stage. These approaches can be classified into two principal categories: traditional methods, which are also referred to as parametric methods and data-driven methods, also known as deep learning methods.

3.4.1 Traditional RFI mitigation techniques

Traditional methods, namely parametric methods in signal processing entail utilisation of estimated signal characteristics to enhance the quality of the recorded signals. These methods require prior knowledge of certain signal parameters, such as frequency content or amplitude modulation. Mathematical models are then applied to distinguish and isolate the target signal from RFI. This approaches particularly advantageous when the signal properties are known or can be accurately estimated, as it allows for targeted and efficient interference mitigation. In the post-correlation stage, Offringa et al. (2010a) provided a comprehensive overview and categorisation of RFI mitigation methods and builds on existing methods to provide new perspectives and powerful tools for automated RFI mitigation strategies.

Regarding the first type of methods, threshold-based RFI mitigation techniques operated under the assumption that RFI exhibits a higher amplitude compared to the astronomical data. Consequently, these methods identify and eliminate strong RFI by establishing specific threshold levels, making them both the most intuitive and the most extensively utilized approaches in the field. In a earlier study, Maslakovic et al. (1996) applied a thresholding technique subsequent to modeling the temporal waveform signal using a discrete wavelet transform.

The widely recognized Cumulative Sum (CUSUM) method, initially introduced by Page (1954) for statistical process control, functions by monitoring the cumulative total of consecutive sample measurements. If this cumulative total exceeds a dynamic threshold, it triggers an alert, indicating the need for adjustments to ensure data quality. The threshold is commonly set based on the probabilities associated with two distribution parameters, enabling a responsive mechanism to variations within the dataset. To adapt the CUSUM method for RFI mitigation, one can utilize the total observed power, or the power received at a specific frequency by an individual antenna, as the sequence of input values for the CUSUM analysis. The initial proposal for this application was made by Baan et al. (2004a), after which Offringa et al. (2010a, 2012) made further improvements, proposing more efficient RFI mitigation algorithms, including VarThreshold, SumThreshold and AOFlagger. Among these methods, the SumThreshold algorithm stands out as the prevalent choice in current radio telescope processing pipelines, attributed to its demonstrated reliability and efficiency (Offringa et al., 2010b; Peck and Fenech, 2013).

The other type of approach assumes that astronomical data signals follow some statistical distribution, usually a normal distribution, and the occurrence of RFI can break this distributional assumption. On the base of this assumption, astronomers use statistical tools to identify RFI (Fridman, 2008; Bhat et al., 2005).

In addition, this assumption is expressed in the time-frequency plane by regarding the function representing the astronomical data signal as having a smooth plane, while the function representing the RFI will have sharp edges. This approach requires first finding a suitable plane, and then calculating the residuals between the data and the assumed plane, which consists of the systematic noise and the RFI signals. Then we set a threshold on the residuals, and those with higher amplitude are labelled as RFI. Winkel et al. (2007) proposed an algorithm that first performs a 2D baseline fit and then identifies and excludes RFI signals through a series of iterative steps. Athreya (2009) capitalized on the distinctive fringe-stopped correlator output behavior in the presence of RFI to eliminate persistent spatial and temporal RFI sources.

When confronting multibeam receiver systems, Kocz et al. (2010) implemented spatial filtering to adeptly discern and excise RFI from temporal signals, employing a Singular Value Decomposition (SVD) of the empirical covariance matrix derived from the Fourier transforms of the input signals. Pen et al. (2009) explored Singular value decomposition (SVD) for RFI identification, while Zhao et al. (2013) utilized techniques with the help of principal component analysis (PCA). In more recent developments, Finlay et al. (2023) has harnessed the predictable trajectories of RFI for a concurrent approach to its elimination and the calibration process.

It is worth mentioning that there are algorithms that have been specifically developed to remove RFI from pulsar observations, taking into account the characteristics of pulsars and combining them with previously mentioned statistical techniques and thresholding algorithms, such as the best known and widely used Coastguard algorithm (Lazarus et al., 2016) and Clfd algorithm (Morello et al., 2019). The PPSD algorithm proposed by Song et al. (2021) identifies the pulsar phase via a TOA model and isolates the pulse profile. It labels transient RFI in off-pulsar data using statistical thresholds, iteratively refining until the data matches white gaussian noise (WGN) distribution, then replaces RFI with WGN. Shan (2023) present a novel M-type estimator using the LnCosh criterion, integrated into iterative shrinkage-thresholding algorithm (ISTA) and Fast ISTA, offering a practical framework for robust RFI excision.

3.4.2 Data-driven RFI mitigation approaches

In recent years, there has been a notable advancement in the field of machine learning and deep learning technologies, which have demonstrated exceptional performance on a multitude of tasks, including computer vision. This has led to a surge in global interest and research activity in this area. Deep learning, a subfield of machine learning, learns complex representations of data by building multi-layer neural network models. These models are capable of automatically extracting features and performing pattern recognition. They are also known as data-driven algorithms because the performance of deep learning algorithms relies heavily on the large amounts of training data that are used to guide the network learning process, thus enabling the algorithms to recognise and predict new and unseen data. In short, the power of deep learning models stems from the inherent laws and patterns they learn from large amounts of data. Applying deep learning techniques to the RFI removal problem, researchers often resort to convolutional neural networks (CNN). In this application, traditional astronomical data in the time-frequency domain is treated as an image, and the advanced capabilities of CNN in image recognition are utilised to identify and remove non-astronomical signals, i.e. RFI, from the data. This approach typically involves converting the two-dimensional plane of the time-frequency data into an image-like format, and then training a convolutional neural network to differentiate and isolate the astronomical signals from the interferences, thus improving data quality and provide clearer observations for subsequent astronomical analyses. Depending on its training method and target task, deep learning can be classified into the following categories: supervised learning, unsupervised learning, and generative adversarial networks (GAN).

Supervised learning in RFI mitigation

Supervised learning is the most common approach in deep learning, which uses labelled training data to train models. In addressing the RFI mitigation problem, a pioneering approach within supervised learning is to conceptualize RFI mitigation as an image segmentation task. The innovative methodology was initially proposed by Akeret et al. (2017a), who perceived the time-frequency domain data as an image and employed image segmentation techniques to delineate and segregate RFI. Utilizing the U-Net convolutional neural network, a widely recognized architecture in the field of image segmentation, this work demonstrated the feasibility of this transformation and showcased its exceptional capability in RFI identification. To enhance model performance, Yang et al. (2020) have integrated residual blocks and batch normalization into their design, resulting in the RFI-Net, tailored for RFI detection in data from the FAST radio telescope. Yan et al. (2021) explored the significance of atrous convolution through the introduction of the AC-UNet model. Gu et al. (2024) proposed EMSCA-UNet model, which employs multi-scale convolutional operations to extract RFI features of different scale sizes, and uses the attention mechanism to assign different weights to the extracted RFI feature maps so that the model can focus on the key features for RFI detection. Meanwhile, the R-Net, proposed by Vafaei Sadr et al. (2020), has exhibited robustness against both simulated and actual data sets, leveraging the power of transfer learning. It should be noted that the aforementioned methods are proposed within a supervised learning context, meaning they require pre-labeled datasets for training. For example, Hamid et al. (2022) utilized predictions from the parametric methods Coastguard and Clfd as the ground truth to train their PSRFINET model for RFI detection in pulsar data. Alternatively, the problem of RFI mitigation can be solved by transforming it into an object detection problem. Ghanney and Ajib (2020) attempt this approach using the YOLOv3 model and compare its performance with that of an unsupervised convolutional auto-encoder.

The aforementioned algorithms utilize deep convolutional neural networks for the identification of regions tainted by RFI, yet they fall short in rectifying the affected astronomical signals. To address this limitation, certain researchers have employed auto-encoder (AE) as a means to restore data compromised by RFI. By stacking several convolutional denoising autocoders, Chang et al. (2023) proposes a denoising autocoder called denoising autoencoder for auroral radio emissions (DAARE) to remove RFI. The concept has also been extended to remove RFI in 5G base stations, as demonstrated by the work of (Owfi and Afghah, 2023). It is worth pointing out that AE is commonly used in unsupervised learning, but these latest discussed works use labeled data for training.

Unsupervised learning in RFI mitigation

Unsupervised learning is a method of learning the structure and patterns of data through the analysis of datasets without the necessity of pre-labelled or classified training data. In unsupervised learning, the algorithmic process is oriented towards understanding the intrinsic structure and relationships of the data, rather than towards predicting a specific output value. The most notable advantage of unsupervised learning over supervised learning is that it does not impose the restriction of requiring labelled data, which is often challenging to obtain in the context of RFI removal problems.

The variational autoencoder (VAE) operates within the realm of unsupervised learning, where it performs dimensionality reduction by employing an encoder to compact the data into a condensed latent form. Subsequently, a decoder is utilized to reconstruct the original data from this latent representation. The VAE is trained to discern the underlying structure of the data by minimizing the discrepancy between the actual input and its reconstructed counterpart. It is frequently applied in various tasks, including the removal of noise from data, the extraction of essential features, and the generation of new data instances. Mesarcik et al. (2020) employed VAE to project the high-dimensional radio spectral data into a low-dimensional latent space, and then demonstrated the performance of the model in different dimensional projections by using a support vector machine (SVM) classifier. In a subsequent work, Mesarcik et al. (2022) introduced an unsupervised technique, the nearest latent neighbours (NLN) algorithm. This method employs the generative adversarial framework to identify and accurately localise the presence of RFI, even in the absence of explicit exposure to RFI. Depending on the fact that the real signals of astronomical signals are thermal and therefore can be regarded as Gaussian random process and that such signals are incompressible, a VAE network is proposed by Saliwanchik and Slosar (2022) to identify and extract the compressible information related to RFI in the data stream, and thus identify and remove RFI. Based on β-VAE, Ma et al. (2023) proposed a fast and modular RFI removal algorithm. Wang et al. (2020) proposed pseudo inverse learning based auto-encoder (PILAE), which removes both broad-band and narrow-band RFI from pulsar data by training a auto-encoder with a single hidden layer while being able to retain most of the pulsar signal.

Moreover GAN have also been successfully implemented in the RFI removal problem. GAN is a deep learning model comprising two principal components: firstly, a generator, which is responsible for the generation of synthetic data that is capable of mimicking genuine samples; and secondly, a discriminator is tasked with distinguishing between authentic and artificial data. During the training phase, the two components engage in a competitive process, driving each other to enhance their performance continuously. Ultimately, the generator is capable of producing high-quality samples that are indistinguishable from genuine data. Vos et al. (2019) trained a generator network for separating astronomical signals from RFI in mixed data. Li et al. (2021) introduced RFI-GAN model founded on the Pix2Pix framework, which re-conceptualizes the RFI detection challenge as an image-to-image translation task. This approach entails the training of a pair of deep neural networks in a competitive environment to produce a binary mask that identifies and neutralizes RFI. Further, by integrating attention module and residual connectivity in the generator module for efficient segmentation, O.M. and R. (2023) proposed attention-segmentation GAN (SegGAN).

3.4.3 Summary and challenges

We have already provided an overview of existing RFI mitigation methods, including traditional signal processing techniques and advanced data-driven approaches.

The traditional RFI mitigation methods are primarily focused on signal processing techniques, which leverage known signal characteristics or statistical properties to identify and eliminate RFI. The benefits of these classical techniques lie in their high computational efficiency, ease of implementation and integration, and stability and reliability under specific conditions. These methods are typically founded upon established theoretical frameworks, possess controllable parameters, and do not necessitate a substantial amount of data, rendering them well-suited for real-time processing and resource-limited environments. The reliance of traditional RFI mitigation methods on prior knowledge of signal characteristics may, in some cases, limit their ability to recognise unknown interference. Moreover, they necessitate the pre-setting of parameters, which can result in a decline in performance in response to a variable RFI environment, and may affect the efficacy of the mitigation process or cause data loss due to the improper setting of parameters. The utilisation of data-driven algorithms, particularly those based on deep learning techniques, as elucidated in Section 3.4.2, presents a multitude of substantial benefits. The primary advantage of these algorithms is that they automatically extract complex features from data through the construction of multi-layer neural network models. This reduces the reliance on prior knowledge required in traditional signal processing techniques and enhances the identification capabilities for novel and unknown RFI sources. Secondly, these algorithms utilise extensive training data to optimise the generalisability of the models, thereby enabling them to perform well not only on the training data set but also to adapt to new observation conditions and RFI environments. Lastly, data-driven algorithms facilitate the acceleration and precision of RFI detection and removal through the implementation of automated processing workflows. This is especially advantageous for the management of extensive radio astronomical datasets, markedly augmenting the overall efficacy and impact of data processing.

While deep learning methods are an effective approach to RFI mitigation, they are not without notable limitations and challenges. Firstly, supervised learning approaches are dependent on the availability of substantial quantities of accurately labelled real data, which is scarce in the domain of radio astronomy. The dearth of such labelled data impedes the training of models capable of effective RFI recognition. In contrast, unsupervised learning frequently fails to achieve the same level of performance as its supervised counterpart, thereby further complicating the development of robust RFI detection algorithms. In the second place, it is common approach to flag RFI locations and then remove these pieces in subsequent processing, which causes the astronomical signals to be lost along with the interference. Despite some studies employing autoencoders to restore astronomical data contaminated by RFI, achieving promising results, there are still several issues that need to be addressed.

Résumé du chapitre

Les interférences radioélectriques (RFI) sont un problème majeur en astronomie radio, qui fait référence aux signaux indésirables qui polluent le spectre des fréquences radio utilisées pour diverses observations astronomiques. Ces interférences proviennent principalement de signaux radio générés par des activités humaines, comme la radio FM, la diffusion télévisuelle, les services téléphoniques mobiles et les systèmes de communication par satellite, notamment pour le positionement (GPS). De plus, les activités terrestres à proximité des observatoires, telles que les systèmes d'éclairage, les lignes électriques et les clôtures électriques, peuvent produire des perturbations transitoires ou persistantes. Les RFI posent des défis significatifs aux télescopes radio pour la détection de signaux astronomiques faibles, ce qui peut entraîner des interprétations erronées des phénomènes astronomiques et des résultats de détection erronés. Une gestion et une atténuation efficaces de RFI sont cruciales pour assurer la fiabilité des observations astronomiques.

Pour répondre aux défis posés par les RFI, le domaine de l'astronomie radio utilise une variété de stratégies et de techniques, classées par étape d'observation. Cela comprend les mesures de prévention dans les observatoires, les stratégies de prédétection, les techniques de pré-décorrélation et les méthodes de post-décorrélation. Au niveau de l'observatoire, la prévention des RFI implique la gestion du spectre, qui réserve des bandes de fréquences nécessaires pour les observations astronomiques et restreint les activités potentiellement perturbatrices dans ces bandes. L'Union Internationale des Télécommunications (ITU) joue un rôle clé en fournissant des recommandations et des directives pour établir des zones sans RFI et pour réguler l'utilisation des fréquences autour des observatoires (RA, 2003).

Les stratégies de pré-détection sont mises en œuvre au niveau du récepteur afin de minimiser les RFI avant qu'elles n'affectent le processus d'acquisition de données. Cela comprend l'installation de filtres passe-bande ou de filtres haute/basse-fréquences pour filtrer les signaux RFI connus et de forte énergie, ainsi que l'arrêt du processus d'acquisition de données pendant les périodes de forte réception de RFI. Les techniques de pré-corrélation s'appliquent avant que les données ne subissent une corrélation, en utilisant le traitement numérique en temps réel, la suppression de bruit adaptative et le filtrage spatial pour réduire les RFI.

Les méthodes de post-corrélation reposent sur des techniques de traitement numérique avancé pour identifier et éliminer les RFI. Cela comprend le filtrage numérique, l'annulation adaptative de RFI à l'aide de canaux ou d'antennes de référence, ainsi que l'application d'algorithmes de seuillage. La phase de post-corrélation est plus coûteuse en temps de calcul et demande plus de ressources, s'appuyant sur le travail fondamental des stratégies de pré-corrélation pour garantir l'exactitude et la fiabilité des données astronomiques radio (Offringa et al., 2010a).

Avec les progrès des technologies d'apprentissage automatique et en particulier d'apprentissage profond, les approches guidées par les données pour l'atténuation de RFI se sont avérées être une promesse significative. Ces méthodes extraient automatiquement des caractéristiques complexes à partir des données à l'aide de modèles basés sur des réseaux neuronaux à plusieurs couches, ce qui réduit la dépendance par rapport aux connaissances préalables et améliore l'identification des sources RFI nouvelles et inconnues. Les modèles d'apprentissage profond, tels que les réseaux neuronaux convolutifs (CNN), traitent les données du domaine temps-fréquence comme des images et utilisent les capacités avancées de reconnaissance d'image pour identifier et éliminer les RFI. Ces méthodes peuvent être classées en fonction de leur technique d'apprentissage : apprentissage supervisé et apprentissage non-supervisé.

Bien que les méthodes d'apprentissage profond soient efficaces, elles font face à des défis, notamment la dépendance à de grandes quantités de données réelles étiquetées de manière précise, qui sont rares en astronomie radio. De plus, bien que certaines études aient utilisé des autoencodeurs pour restaurer les données compromises par des RFI, plusieurs problèmes doivent encore être abordés.

Part II

Contribution to RFI mitigation

Chapter 4

Framework for simulating RFI corrupted dynamic spectra observation

4.1 Introduction and motivation

As discussed in Section 3.4.3, the effectiveness of supervised deep learning techniques is inherently contingent upon the utilisation of datasets that have been annotated with precise labels for both training and validation purposes. Furthermore, even in the context of unsupervised learning algorithms, the utilisation of labelled datasets is essential for validating their performance in RFI detection.

In the past, astronomers could obtain labeled datasets through manual annotation. However, the current era of advanced radio telescopes has introduced an insurmountable challenge to this approach, given the enormous volume of data they produce. Furthermore, the objective of RFI mitigation is not confined to the removal of interference; it also involves the recovery of the original astronomical signals to prevent data loss that is intrinsic to methods that focus solely on identifying and eliminating RFI. This necessitates the availability of paired datasets that include both clean astronomical data and counterparts affected by RFI, a requirement beyond the scope of traditional labeling techniques. While prior research has explored comparable strategies for data generation, none have been tailored explicitly for the purpose of pulsar observations Akeret et al. (2017b); Asad et al. (2021); DeBoer et al. (2017). Therefore, the development of a framework that can simulate pulsar observation data is essential to enable the generation of such paired datasets.

By leveraging the pulsar observation data recorded by the NenuFAR radio telescope, we have established a data generation framework that simulates NenuFAR like observational data. The framework has been designed to produce datasets that are accurately labelled, which will be instrumental in the subsequent development of supervised learning algorithms capable of restoring pulsar observation data that has been compromised by RFI.

4.2 Observation model

The four data modes after the correlation provided by NenuFAR for pulsar observations have been described in Section 2.3. The present study seeks to remove RFI and recover pulsar signals from the contaminated observations in the dynamic spectra mode, which essentially involves calculating the magnitude of the discrete Fourier transform (DFT) of the raw data.

The rationale for simulating data in the dynamic spectrum mode can be distilled into three principal points. Firstly, the resolution of the dynamic spectrum, represented on a time-frequency plane through a discrete windowed analysis, is directly determined by the parameters of the DFT. This approach allows for the modification of DFT parameters in order to align them with the specific temporal and spectral resolution requirements of our study. Secondly, in comparison to pure time-domain analysis, RFI characteristics are more pronounced in the time-frequency plane, thereby facilitating the identification of weak RFI signals that may not be evident in single time or frequency analyses. Finally, while numerous studies have focused on RFI removal in pulsar observations, particularly in the folded mode, there is a relative dearth of research in the dynamic spectrum mode. Therefore, the development of an effective RFI mitigation algorithm for this mode addresses an existing research gap.

Our objective is to engineer a data generation system capable of producing paired datasets consisting of both pulsar signals and their counterparts afflicted with RFI. To achieve this, we will independently generate both the RFI and pulsar signals and then combine them to produce a RFI-corrupted spectrum. This approach assumes that the signals are additive when represented in the time-frequency plane (Fridman and Baan, 2001), enabling an accurate simulation of their interplay under realistic observational conditions. On the time-frequency plane, the recorded signals can be represented by a function $S(n, k)$, where n indexes the frequency channel and k denotes the temporal bin. The signal model is hypothesized to be decomposable into its constituent parts, represented mathematically as

$$S(n, k) = P(n, k) + R(n, k) + E(n, k) \tag{4.1}$$

where $P(n, k)$ is the pulsar signal, $R(n, k)$ is a possible RFI component and $E(n, k)$ stands for the system noise and any mismodeling, for $n \in \{1, \ldots, N\}$ (n is the frequency index and N is the number of spectral bins) and $k \in \{1, \ldots, K\}$ (k is the time index and K is the number of temporal bins).

A depiction of a dynamic spectrum, generated in accordance with the model presented in Equation (4.1) is provided in Figure 4.1. Within this illustration, the pulsar signals $P(n, k)$ is characterized by exponentially decaying curves, while RFI signals $R(n, k)$ manifests as vertical or horizontal lines and small clusters of dots.

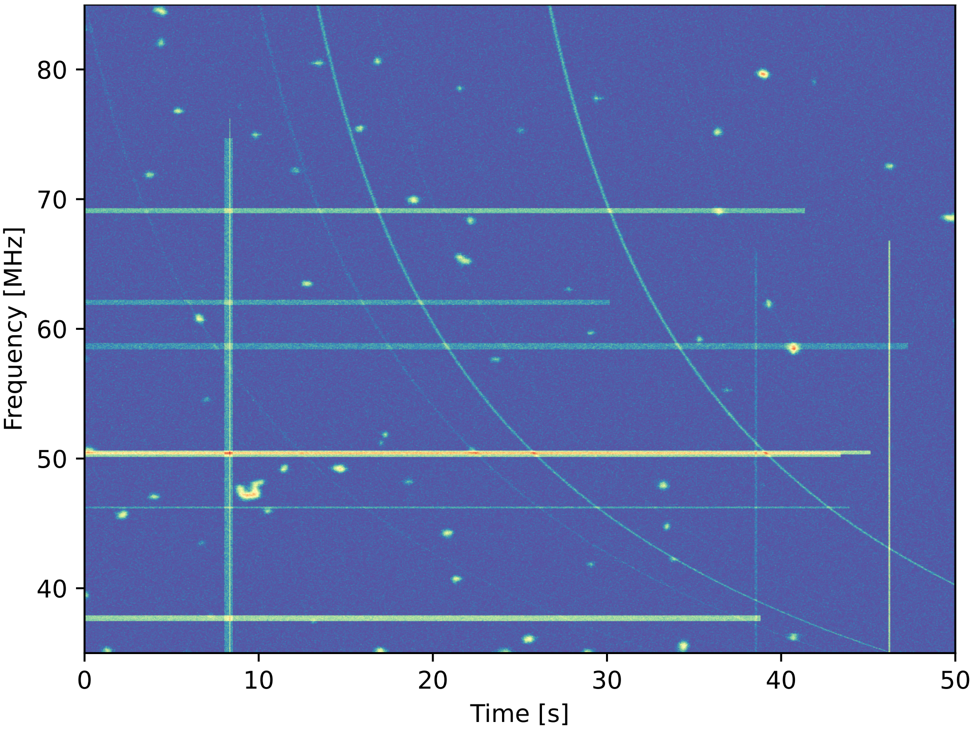

Figure 4.1: An illustration of the dynamic spectrum generated on the basis of the proposed model delineated in Equation (4.1).

4.3 Simulation of pulsar signal

Pulsar signals, characterized by their distinctive periodicity and highly stable rotational periods, serve as ideal candidates for astronomical timing studies. In an effort to accurately simulate these signals, this work has developed a simulation approach have been developed that includes the fundamental characteristics of pulsar emissions as introduced in Section 1.3. This approach leverages the diverse properties of pulsars to generate signals that closely mimic the observed phenomena.

Time-frequency representation	
N	Number of spectral bins (channels)
n	Index of the spectral bin (channel)
K	Number of temporal bins
k	Index of the temporal bin
δt	Temporal resolution
δf	Spectral resolution
$S(n,k)$	Dynamic spectrum (full signal)
$P(n,k)$	Pulsar signal
$R(n,k)$	RFI signals
$E(n,k)$	System noise

Pulsar modeling	
D	Number of periods in the observation window
d	Index of the period in the observation window
ρ	Pulsar period
$A(n,k)$	Integrated profile
$\gamma_{\mathrm{P}}(\cdot;\sigma^2)$	1D Gaussian kernel
SNR_d	Signal-to-noise ratio in the dth period
τ_n	Dispersion delay
DM	Dispersion measure
L	Number of Gaussian kernels
a_ℓ	Amplitude of the ℓth Gaussian kernel
σ_ℓ	Width of the ℓth Gaussian kernel
μ_ℓ	Location of the ℓth Gaussian kernel

RFI modeling	
J	Number of RFI
SNR_j	Signal-to-noise ratio of the jth RFI
$M_j(n,k)$	Binary mask to locate the jth RFI
$\gamma_{\mathrm{R}}(\cdot,\cdot;\sigma_{\mathrm{T}}^2,\sigma_{\mathrm{F}}^2)$	Separable 2D Gaussian kernel
α	Probability of occurrence of nbct and bbt RFI
β	Granularity parameter to generate nbt RFI

Table 4.1: Notations used to describe the RFI-corrupted dynamic spectra generated by the simulation protocol.

4.3.1 Template-Based Model Enhancement

The simulation begins with a template model representing the integrated pulsar profile, which acts as a distinctive 'fingerprint' for the pulsar. The pulsar signal, denoted as $P(n, k)$, features a unique profile represented by $A(n, k)$. The simulation of pulsar signal involves examination their pulse profiles,which form the basis of replicating the temporal structure of the pulsar signals.

The pulse profile of a pulsar signal is generally stable in terms of morphology during observations at a given frequency, aside from variations in energy (Lorimer and Kramer, 2004). Consequently, in the simulation of pulsar signals, such temporal variations are neglected, under the assumption that the shape of the pulse profile remains constant over time. Instead, changes in energy, or equivalently, amplitude, are considered, meaning that the pulse profile for each cycle may have a different amplitude. Therefore, the pulsar signal over D periods can be expressed as:

$$P(n,k) = \sum_{d=1}^{D} \mathrm{SNR}_d \times A(n, k - \tau_n - d\rho) \tag{4.2}$$

where SNR_d represents the the adjustment factor for the signal-to-noise ratio specific to the d-th observation period. ρ denotes the rotational period of the pulsar, defining the interval between successive pulses. τ_n is the time delay that varies with frequency, an effect attributed to the dispersion of the radio waves, which will be detailed in Section 4.3.2

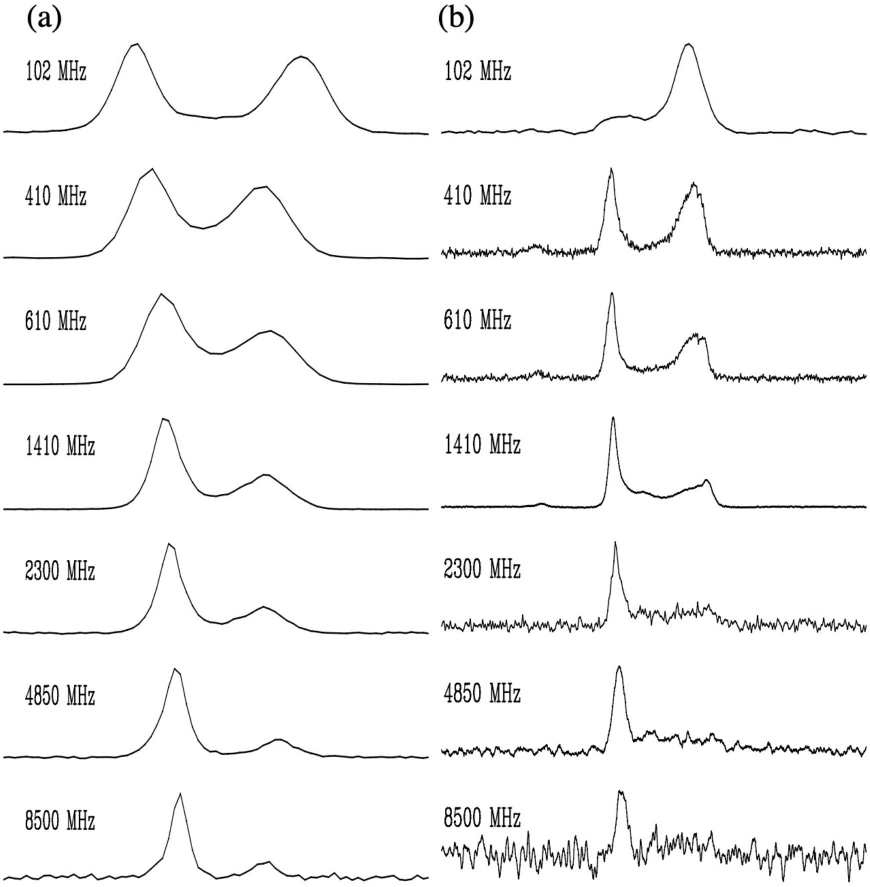

Figure 4.2: Example of pulse shape's evolution with frequency. a) The 1.16s pulsar B1133+16; b) The 16ms pulsar J2145-0750 (Lorimer and Kramer, 2004).

As shown in the figure 4.2, with changes in the observation frequency, the pulse profile of pulsars undergoes significant alterations. Specifically, at lower frequencies, most pulsars exhibit a systematic increase in pulse width, accompanied by a separation in the pulse profile, a phenomenon known as radius-to-frequency mapping (Lorimer and Kramer, 2004). The emergence of this phenomenon renders the modeling of pulse profiles across different frequencies nearly impossible, as the patterns of frequency-dependent variation can differ significantly among pulsars. The diversity in pulsar magnetic field configurations, the pronounced impacts of dispersion and scattering at lower frequencies, and the unique physical characteristics of each pulsar are all factors that contribute to the complexity of pulse profile variations. (to be confirmed)

Despite the acknowledged complexity in modeling pulse profiles across a wide range of frequencies due to the variability among pulsars, it is important to note that for the specific frequency band of interest in our study, the variations are minimal. Specifically, within the 10 MHz to 85 MHz range covered by the NenuFAR telescope, the changes in pulse profiles are negligible. For clarity, Fig. 4.3 presents a visual comparison of the pulse profiles for the pulsar B1919+21, as captured by the NenuFAR telescope. This is attributed to the relatively small frequency span, where the effects of dispersion and scattering, as well as the influence of magnetic field configurations, are less pronounced compared to larger frequency intervals. In light of this, we make the simplifying assumption that the pulse profiles remain constant across the frequencies within our observation range. This assumption allows us to bypass the intricate modeling challenges associated with large frequency sweeps while maintaining the fidelity of our simulations within the operational bandwidth of the the NenuFAR observatory.

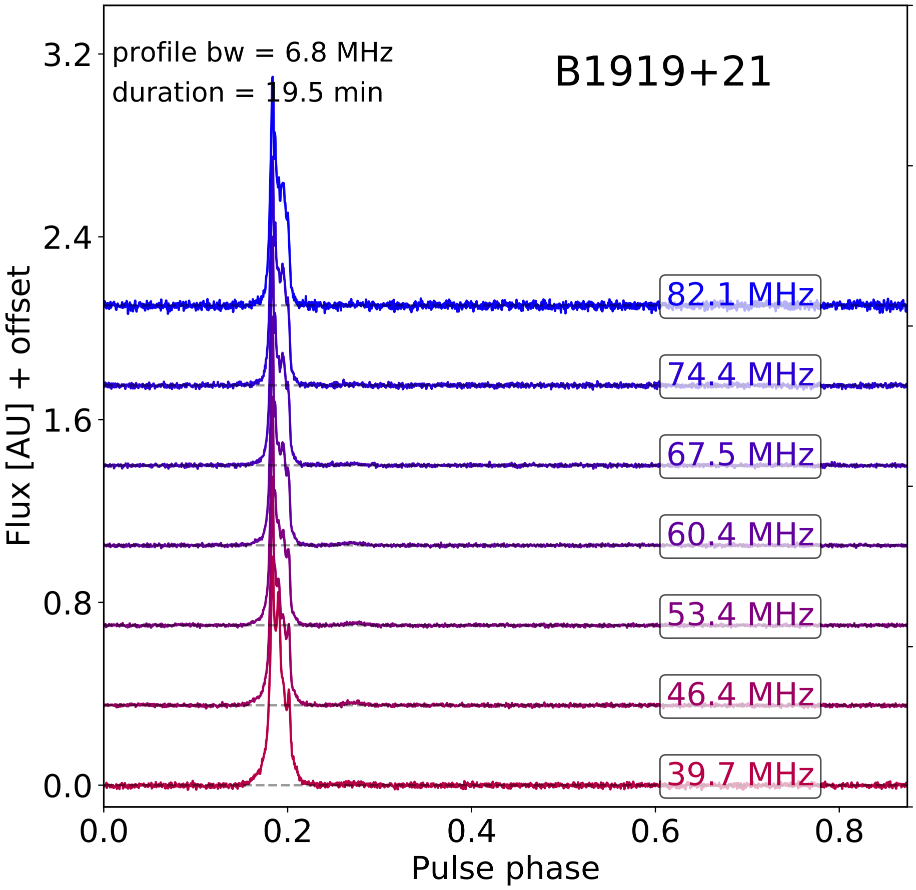

Figure 4.3: Pulse profiles of the pulsar B1919+21, which were recorded over a range of frequencies by NenuFAR telescope. The observation period was integrated for a total duration of 19.5 minutes.

To ensure broad applicability of the methods presented in this manuscript, the generalizing assumption made is that the pulsar profile is entirely characterized by a consistent template across all frequencies. This template is formulated as a weighted linear combination of L Gaussian-shaped components, allowing for a versatile representation of the pulse profile. The equation for the pulsar profile $A(n, k)$ is expressed as:

$$A(n,k) = \sum_{\ell=1}^{L} a_\ell \gamma_{\mathrm{P}}(k - \mu_\ell; \sigma_\ell^2) \tag{4.3}$$

Here, L denotes the total number of Gaussians used to construct the profile. Each component is uniquely characterized by its index ℓ, which serves to distinguish and enumerate the individual Gaussian elements contributing to the overall pulse shape. The amplitude α_ℓ, mean μ_ℓ, and variance σ_ℓ^2 are associated to the ℓ-th component. The Gaussian function γ_{P} is defined as:

$$\gamma_{\mathrm{P}}(t;\sigma^2) = \exp\left(-\frac{t^2}{2\sigma^2}\right) \tag{4.4}$$

4.3.2 Dispersion measurement simulation

As radio pulses traverse the interstellar medium and, in some cases, the intergalactic medium, they encounter a dispersion delay that affects the signal's temporal structure. To address this in our simulation, we incorporate a dispersion measurement that reflects the delay experienced by the pulses due to the electron column density, known as the dispersion measure (DM).

The dynamic spectrum mode of our radio telescope is capable of performing a de-dispersion process, which corrects for the dispersion effects across the frequency spectrum. To ensure our simulated pulsar signals are representative of this phenomenon, we have designed our model to include dispersion effects. Specifically, the dispersion delay for each frequency bin n is calculated using the formula:

$$\tau_n = 4.15 \times \frac{\mathrm{DM}}{n^2 \delta f^2 \delta t} \tag{4.5}$$

In this equation, δf and δt are the spectral and temporal resolutions of the dynamic spectrum. The constant 4.15 is derived from the speed of light and the electron rest mass, and it scales the DM to the observed delay.

This adjustment to the dispersion measurement ensures that our simulated signals accurately mimic the dispersion effects observed in real pulsar data, thereby enhancing the realism and applicability of our simulation model.

4.4 Simulation of RFI signal

In the previous Section 3.1, the definition of RFI was introduced and its potential sources were explored. Building upon this foundation, we referenced the extensive work of (An et al., 2017; Baan, 2011), who provided a detailed characterization of prevalent RFI types, illustrating their spectral signatures through comprehensive spectral analysis. Further to this foundation, this work incorporates actual observational data obtained from the NenuFAR telescope. By examining the discernible traits present in the dynamic spectra, I have identified and categorized RFI by its shape in the frequency and the time time into three distinct classes: *i)* Narrow-band transient (nbt) RFI, *ii)* Narrow-band continuous (nbct) RFI, *iii)* Broad-band transient (bbt) RFI.

To address the variability observed in the spatial and spectral characteristics of RFI, this work employs a methodical approach that involves the decomposition of the composite RFI signal, represented as $R(n,k)$, into a sum of J individual RFI components. Each component is defined by a unique two-dimensional Gaussian template, expressed by $\gamma_R(t,f;\sigma_{\mathrm{T}}^2,\sigma_{\mathrm{F}}^2) \triangleq \exp\left(-\frac{t^2}{2\sigma_{\mathrm{T}}^2}\right)\exp\left(-\frac{f^2}{2\sigma_{\mathrm{F}}^2}\right)$, with variances σ_{T}^2 and σ_{F}^2 that respectively dictate the temporal and spectral extent of the RFI pattern. It is assumed that RFI does not undergo dispersion, thus the dispersion measure (DM) is set to zero DM $= 0\mathrm{pc\ cm}^{-3}$.

The formulation of the overall RFI signal $R(n,k)$ is mathematically represented as:

$$R(n,k) = \sum_{j=1}^{J} \mathrm{SNR}_j \times \gamma_R(n - n_j, k - k_j; \sigma_{n_j}^2, \sigma_{k_j}^2) \tag{4.6}$$

where SNR_j denotes the signal-to-noise ratio of the j-th RFI instance, scaling its power, while n_j and k_j identify the spectral and temporal central positions of the RFI. The spreads in the spectral and temporal axes are determined by $\sigma_{n_j}^2$ and $\sigma_{k_j}^2$, respectively, enabling a detailed portrayal of each RFI component's spatial and spectral attributes.

To introduce randomness in the placement of RFI across the dynamic spectrum, the model presented in the Equation (4.6) is reformulated to incorporate binary masks. These masks are generated randomly, embodying specific statistical traits that reflect the diverse time-frequency patterns of RFI. The reformulation is expressed as:

$$R(n,k) = \sum_{j=1}^{J} \mathrm{SNR}_j \times M_j(n,k) * \gamma_R(n,k;\sigma_{n_j}^2,\sigma_{n_k}^2) \tag{4.7}$$

where the symbol $*$ represents the two-dimensional convolution operation, which disperses the RFI signal across the spectrum according to the Gaussian function γ_R.

The binary mask $M_j(n,k)$ is defined by the following equation:

$$M_j(n,k) = \delta(n-n_j, k-k_j) = \begin{cases} 1, & \text{if } n = k_j \text{ and } k = n_j \\ 0, & \text{otherwise} \end{cases} \tag{4.8}$$

This mask assumes the value of 1 when an RFI event is centered at the frequency bin n_j and the time instant k_j, and 0 in the absence of RFI at those coordinates. Two distinct methodologies are applied for the random generation of these masks, corresponding to different classes of RFI. The procedures for generating these masks are detailed in the subsequent discussion.

4.4.1 Extended RFI in time and frequency

Narrow-band continuous-time RFI and broad-band transient RFI are two prevalent types of interference encountered in actual radio astronomical observations. A case in point is the NenuFAR telescope, where the frequency band of 36-37MHz is often affected by broad-band transient RFI from human activities. Given the general independence of such RFI instances, the generation of their corresponding masks $M_j(n,k)$ can be effectively modeled using a Bernoulli distribution with probability α as shown below:

$$\mathbb{P}[M_j(n,k) = \epsilon] = \alpha^{\epsilon}(1-\alpha)^{1-\epsilon} \tag{4.9}$$

Here, $\epsilon \in \{0,1\}$, indicating the presence or absence of RFI. The parameter α represents the expected value $\mathbb{E}[M(n,k)]$, which corresponds to the probability of RFI occurrence in the dynamic spectrum. Consequently, this parameter is a pivotal parameter that regulates the mean count αNK of spectrally and temporally extended RFI instances, categorized as nbct RFI and bbt RFI.

4.4.2 Narrow-band transient RFI

In order to simulate the clustered impact of narrow-band transient (nbt) RFI on dynamic spectrum, the generation of mask entries cannot rely on independent processes. Instead, a Markov random field (MRF), known for introducing structural correlations across the spectrum, is employed. This MRF is formulated as a multilevel logistic model, often referred to as an Ising model, and is mathematically expressed as follows:

$$\mathbb{P}\left[M_j(n,k) = \epsilon \mid \mathcal{M}_j(n,k)\right] \propto \alpha^{\epsilon}(1-\alpha)^{1-\epsilon} \exp\left[\beta \sum_{m \in \mathcal{M}_j(n,k)} \delta\left(m-\epsilon\right)\right] \tag{4.10}$$

where $\epsilon \in \{0,1\}$, with $\epsilon = 1$ indicating the presence of RFI and $\epsilon = 0$ indicating its absence. $\mathcal{M}_j(n,k)$ represents the neighboring set of the mask entry $\mathcal{M}_j(n,k)$ in the time-frequency plane, structured according to a 4-order neighborhood system:

$$\mathcal{M}_j(n,k) = \{M_j(n-1,k), M_j(n+1,k),\ M_j(n,k-1), M_j(n-1,k)\} \tag{4.11}$$

The granularity parameter β in the model adjusts the level of correlation between adjacent mask entries. When $\beta = 0$, the model reverts to the independent Bernoulli distribution used for modeling spectrally and temporally extended RFI. The simulation of masks using this MRF can be efficiently performed using Gibbs sampling techniques, as detailed in (Li, 2009, Chapter 2).

4.5 Generation of simulated data sets

Following the comprehensive introduction of the methods for generating pulsar signals and various types of RFI signals, this section will now focus on the concrete construction of simulated datasets for training purposes. The specific steps involved in dataset generation will be presented, encompassing parameter configuration, simulation of different scenarios, and the organization and formatting of the final datasets.

4.5.1 Simulated data generation process

A decomposed approach has been adopted to construct our dynamic spectrum model (4.1), allowing for the independent generation of pulsar signals and various types of RFI signals. This method enhances the flexibility of the generation process and significantly enriches the authenticity and diversity of the simulated data.

Initially, a series of pulsar signals were generated based on the characteristics of pulsar emissions, using specific physical and statistical model which has been introduced in the Section 4.3, and stored them in the pulsar database $\mathcal{P}$. Concurrently, we have established a comprehensive RFI signals database, denoted as $\mathcal{R}$, which encapsulates three distinct categories of RFI signals. Each subset is meticulously curated to represent a specific type of RFI: nbt RFI signals are categorized under $\mathcal{R}^{\text{nbt}}$, nbct RFI signals are placed in $\mathcal{R}^{\text{nbct}}$, bbt RFI signals are designated as $\mathcal{R}^{\text{bbt}}$. Subsequently, a pulsar signal was randomly selected from the pulsar database $\mathcal{P}$ and RFI signals from the RFI database $\mathcal{R}$. This approach enables us to simulate a variety of scenarios that might be encountered in actual observations, thereby generating a series of simulated dynamic spectra $S(n, k)$.

Finally, the independent pulsar and RFI signals were combined according to the equation of our dynamic spectrum model (4.1) to form a complete set of simulated data $\mathcal{S}$. The main advantage of this strategy is the significant reduction in required computational power. Specifically, if $|\mathcal{P}|$ and $|\mathcal{R}|$ represent the sizes of the pulsar and RFI databases, respectively, then by combining these two databases, $|\mathcal{S}| = |\mathcal{R}| \times |\mathcal{P}|$ distinct dynamic spectra can be generated. In practice, only $|\mathcal{R}| + |\mathcal{P}|$ independent computations are required, which is much less than $|\mathcal{S}|$, thus achieving a notable increase in computational efficiency.

4.5.2 Simulation parameter configuration

Accurate simulation of observational data necessitates a meticulous configuration of simulation parameters. This subsection details the process of selecting and setting the parameters that govern the characteristics of our simulated datasets.

Global parameters definition

The simulation begins with configuring global parameters, which are essential for establishing the foundation of the simulation environment. These parameters are critical for establishing the structure and resolving power of the dynamic spectra we simulate. Specifically, they include the temporal and frequency resolutions as well as the number of bins in our model.

The selection of parameters for our dynamic spectra simulation is tailored to align with the principal features of signals as observed by the NenuFAR telescope. The observable frequency range of NenuFAR spans from 15 MHz to 85 MHz. However, it is noted that the frequency band below 30 MHz is often saturated with RFI throughout the day. To address this, our simulation focuses on a refined frequency range from 35 MHz to 85 MHz, which is less affected by pervasive RFI and still within the operational bandwidth of the telescope. Within this selected frequency range, the simulation is designed with a granularity that reflects the detailed observations possible with NenuFAR. Specifically, the frequency axis is discretized into $N = 1024$ points, allowing for a comprehensive coverage of the chosen spectrum. This discretization translates to a spectral resolution of $\delta f = 48.828\text{kHz}$. Besides, the time resolution of the dynamic spectra is set to $\delta t = 0.05\text{s}$. Along the time axis, an equal number of 1024 points are utilized, paralleling the configuration of the frequency domain to create a symmetric and high-resolution data grid suitable for the analysis of dynamic phenomena.

Pulsar simulation parameters

The pulsar signal, denoted as $P(n, k)$, is intricately characterized by its integrated profile $A(n, k)$, as detailed in Equation (4.2). This profile is shaped by a template function $\gamma_{\mathrm{P}}(\cdot, \cdot)$, defined in Equation (4.3), which serves as a fundamental representation of unique signature characteristic of pulsars.

To simulate the periodic nature of pulsar emissions, the signal is periodized over the observation interval, with the pulsar period ρ and dispersion measure (DM) parameters being pivotal in this process. These parameters are assigned values that are uniformly and randomly selected within predefined ranges, designed to emulate the variability observed in genuine pulsar signals. The signal-to-noise ratio (SNR) across different periods of the pulsar signal, represented as SNR_d for $d = 1, \dots, D$, is strategically chosen following a log-uniform distribution. This selection reflects the natural logarithmic distribution of pulsar signal strengths and ensures a realistic simulation of the pulsar power levels over D observation periods.

The template function γ_{P} is constructed from up to L Gaussian components, each characterized by its amplitude, localization μ_ℓ, and width σ_ℓ^2. The number of Gaussian components, L does not exceed two, allowing for a concise yet comprehensive representation of the pulsar profile. The specific values for μ_ℓ, and σ_ℓ^2 are uniformly selected from pre-defined sets, ensuring a diverse range of profile shapes that capture the essence of different pulsar emissions.

A comprehensive summary of the permissible ranges for these parameters is presented in Table 4.2, providing a clear reference for the parameter settings used in our simulation study.

Parameter	**Notation**	**Value range**
Period	ρ	$(20, 40)$ [bins]
Dispersion	DM	$(10, 40)$ [pc cm^{-3}]
Number of components	L	$\{1, 2\}$
Amplitude	a_ℓ	$(0.2, 1)$
Localization	μ_ℓ	$(0, \rho)$
Width	σ_ℓ^2	$(0.01, 0.04)$
Power	SNR_d	$(0.01, 20)$

Table 4.2: Parameters of the simulation associated with the pulsar signal.

RFI simulation parameters

The simulation of RFI include three different categories: narrow-band transient (nbt) RFI, broad-band transient (bbt) RFI, and narrow-band continuous (nbct) RFI. Each type is generated using a specific protocol, ensuring a realistic representation of the interference encountered in radio astronomical observations. For bbt RFI and nbct RFI, binary masks are created using a Bernoulli distribution, which allows for the stochastic generation of RFI. In contrast, nbt RFI is generated through a Markov random field, introducing a structured correlation among RFI occurrences, simulating the clustered nature of these transients.

Once the binary masks are established, they are convoluted with a two-dimensional separable Gaussian kernel. This kernel is characterized by its spectral spread σ_{F}^2 and temporal spread σ_{T}^2, parameters that are crucial in shaping the frequency and time-domain profiles of RFI. These spreads are randomly selected from predefined ranges, tailored to reflect the unique attributes of each RFI type. The selection of σ_{F}^2 and σ_{T}^2 is paramount, as they dictate the extent and shape of the RFI impact on the dynamic spectra. The parameters governing the statistical properties of the binary masks, along with the Gaussian kernel parameters, are meticulously detailed in Table 4.3. This table serves as a reference for the parameter ranges that are permissible for simulating each category of RFI, ensuring that the generated interference are both diverse and representative of real-world observations.

Dataset composition

Our simulation strategy systematically integrates pulsar signals from dataset $\mathcal{P}$ and a multitude of RFI signals from dataset $\mathcal{R}$ to create dynamic spectra. The RFI dataset $\mathcal{R}$ is meticulously divided into three subsets, each corresponding to a different type of RFI and ensuring an equal representation of each RFI

Type of RFI	SNR_j	σ_F^2	σ_T^2	α	β
bbt	$(1, 10)$	$(600, 1024)$	$(1, 10)$	$(0, 0.01)$	N/A
nbct	$(1, 10)$	$(1, 10)$	$(600, 1024)$	$(0, 0.01)$	N/A
nbt	$(0, 1)$	$(1, 11)$	$(1, 11)$	0.8	40

Table 4.3: Parameters of the simulation associated with the RFI signal.

category. Through this simulation protocol, we generate three different datasets to form a complete dataset for subsequent training and validation: *i*) the training set, *ii*) the validation set,*iii*) the testing set. The training and validation datasets are constructed from identical pulsar and RFI databases, with the sole distinction being the size of the datasets, denoted as $\mathcal{S}$. This ensures that the models are trained and validated on a consistent range of signal characteristics, allowing for a robust assessment of their performance. In contrast, the testing set is derived from a separate set of pulsar and RFI signals, introducing variability that is crucial for evaluating the generalizability of the trained models to new data.

The specific sizes of these datasets are delineated in Table 4.4, providing a clear overview of the quantities involved in our simulation. This detailed breakdown is essential for understanding the scale of our simulation efforts and the comprehensiveness of the data used to train, validate, and test our models.

	Size of $\mathcal{P}$	**Size of $\mathcal{R}$**	**Size of $\mathcal{S}$**
Training set	20	300	1800
Validation set	20	300	200
Testing set	10	60	200

Table 4.4: Size of the generated training, validation and testing sets.

4.5.3 Simulation scenario construction

In the observation model 4.1, two components are the pulsar signal P and the RFI signal R, which detail how the simulation should be conducted. The final component, E, represents the system noise and any mismodeling elements that have not yet been elaborated upon. Regarding the noise component, two scenarios have been considered to examine the robustness of our proposed method against noise interference. In Scenario 1 (shortened as S_1 hereafter), E is assumed to be zero, indicating an absence of system noise. This assumption is made to simulate an ideal observation condition. In Scenario 2 (shortened as S_2 hereafter), the noise E is modeled as additive white Gaussian noise. Introducing this level of noise, allowing us to assess the method's resilience in the presence of such noise. By contrasting these two scenarios, we aim to evaluate the effectiveness of our method under conditions without noise and with noise.

Furthermore, to explore the ability of the model to handle various types of RFI, four specific cases have been established within each scenario, each distinguished by a different composition of RFI within the set $\mathcal{R}$. Denoted as C_A to C_D, these cases, defined as follow, facilitate to evaluate the perfoamance of the model in dealing with distinct RFI configurations.

- C_A: nbt RFI (pulse-like RFI)
- C_B: nbt RFI + nbct RFI (narrow-band RFI)
- C_C: nbt RFI + bbt RFI (transient RFI)
- C_D: nbt RFI + nbct RFI + bbt RFI (all RFI types)

In summary, our analysis encompasses a comprehensive set of eight simulation scenarios, represented as $\mathsf{S}_\square\mathsf{C}_\triangle$. Here, $\triangle$ takes the values of 1 or 2, indicating the presence or absence of system noise, respectively. Meanwhile, $\square$ represents one of four possible RFI configurations, labeled $\mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D}$. These scenarios are crafted to systematically examine the model's response to different RFI conditions affecting the pulsar signal. For clarity, Figure 4.4 presents a visual representation of the dynamic spectra generated for each of the four cases within scenario S_2.

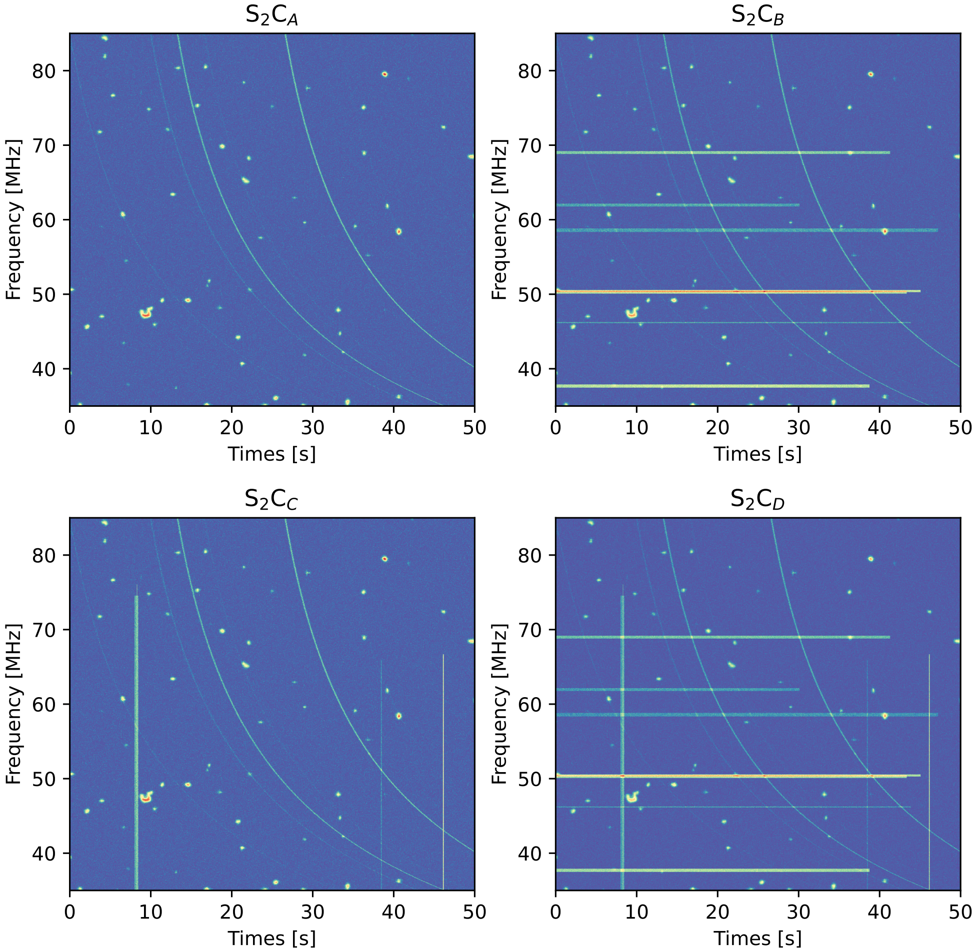

Figure 4.4: Examples of the dynamic spectrum generated according to the proposed protocol for the 4 cases of S_2.

Résumé du chapitre

Pour générer des données semblables à celles fournies par les télescopes radiomodernes, nous avons proposé un cadre de simulation pour simuler des observations de pulsars affectées par des RFI. Ce cadre est développé afin de permettre l'entraînement d'algorithmes d'apprentissage supervisé, capables de restaurer les données de pulsars altérées par les RFI tout en préservant les signaux astronomiques initiaux, évitant ainsi une perte de données.

Nous avons opté pour la simulation de données en mode spectre dynamique en effectuant la transformée de Fourier discrète sur les données brutes. Cette méthode a été choisie en raison de sa souplesse en matière de résolution temporelle et fréquentielle, et de sa capacité à rendre les caractéristiques des RFI plus apparentes, ce qui facilite l'identification des RFI de faible énergie qui pourraient passer inaperçues lors d'analyses temporelles ou fréquentielles séparées. Notre objectif est de concevoir un algorithme efficace pour l'atténuation des RFI dans les spectres dynamiques.

Pour simuler les signaux des pulsars, nous avons employé une méthode qui tient compte de leurs caractéristiques fondamentales, notamment leur périodicité et leurs périodes de rotation stables. Un noyau gaussien a été utilisé pour simuler la structure temporelle des pulsars, en tenant compte du rapport signal-sur-bruit et des effets de dispersion, garantissant ainsi la fidélité des signaux simulés.

Pour simuler les signaux RFI, nous avons classifié les RFI en trois catégories basées sur des données d'observation du télescope NenuFAR : RFI transitoires à bande étroite (nbt), RFI continues à bande étroite (nbct) et RFI transitoires à bande large (bbt). Chaque composante RFI a été représentée par un modèle gaussien bidimensionnel, et des masques binaires et des processus aléatoires ont été utilisés pour simuler la répartition et l'agencement des RFI dans le spectre dynamique.

Dans le cadre de la création de jeux de données simulés, nous avons généré de manière indépendante les signaux des pulsars et les signaux RFI, puis les avons combinés pour créer des ensembles de données simulés complets. Cette démarche a considérablement réduit le coût computationnel. Nous avons également pris en compte l'influence du bruit du système en simulant des scénarios sous différentes conditions de bruit pour évaluer l'efficacité et la stabilité de notre méthode. Ces ensembles de données simulés sont essentiels pour l'entraînement et la validation d'algorithmes d'atténuation des RFI pour les données d'observation des pulsars. Nous avons créé des ensembles d'entraînement, de validation et de test pour nous assurer que les modèles sont entraînés et validés sur une gamme cohérente de caractéristiques de signal, et pour évaluer la capacité du modèle à s'adapter à de nouvelles données. Grâce à ces ensembles de données, nous sommes en mesure d'évaluer de manière globale les performances du modèle dans diverses conditions de RFI.

Chapter 5

Introduction to deep learning and convolutional neural networks

5.1 Introduction to deep learning

Deep learning, a critical subdivision of artificial intelligence, has revolutionized various fields with its ability to discern complex patterns from extensive datasets. The foundational principle of deep learning lies in artificial neural networks, which, through layered structures, echo the sophistication of the human brain. This design enables deep neural networks to excel in tasks that require intricate feature recognition and pattern extraction.

The evolution of deep learning is a testament to the confluence of three pivotal elements: expansive datasets that provide the raw material for learning, leaps in computational power that enable processing these datasets, and the maturation of algorithms that facilitate effective training of these complex networks. These advancements have collectively propelled deep learning to unprecedented heights in areas ranging from image and speech recognition to natural language understanding and beyond.

The essence of deep learning is encapsulated in its iterative learning process, where models are continually refined through exposure to data. This process depends on optimising the model parameters to minimise the prediction error, quantified by a cost function. The cost function plays a pivotal role by calculating the difference between the model's predicted outputs and the actual targets, providing a measure of performance that guides the learning algorithm in adjusting the network's weights.

The journey of deep learning begins with the concept of artificial neural networks, which were inspired by the biological neural networks of the human brain. The perceptron, introduced in the 1950s, laid the groundwork for neural networks. It wasn't until the 1980s and 1990s that the foundations for deep learning were laid with the development of back propagation algorithms and the introduction of more complex network architectures. The 2010s witnessed a surge in the impact of deep learning, particularly with the advent of deep convolutional neural networks that excelled in image recognition tasks. This success was fueled by the availability of big data and the computational power to process it, marking a significant milestone in the field of artificial intelligence.

In essence, deep learning encapsulates a dynamic process of model refinement through the lens of data-driven optimization, aimed at minimizing the discrepancy between predictions and reality as measured by a cost function.

5.2 Foundations of neural networks

A neural network is composed of interconnected nodes, or neurons, which process information through a series of weighted inputs and activation functions. Each neuron receives input from other neurons, processes this input through a mathematical function, and then transmits the resulting output to subsequent neurons. The fundamental unit of a neural network is the neuron, which can be represented by a simple mathematical model. This model comprises a set of weights that determine the influence of each input on the neuron output, as well as a bias term that adjusts the threshold for activation.

5.2.1 Perceptron

The perceptron, introduced by Rosenblatt (1958), represents the most elementary form of a neural network, serving as a foundational linear binary classifier. Inspired by the biological neural networks within the human brain, the perceptron operates on a straightforward principle that is elucidated in Figure 5.1. This schematic representation illustrates the mechanism of perceptron action. Mathematically, the perceptron is encapsulated by the following expression:

$$z = f(\mathbf{x}) = \mathtt{sign}(\mathbf{w} \cdot \mathbf{x} + b) \tag{5.1}$$

where $\mathbf{x}$ denotes the input vector, the weight parameter is $\mathbf{w} \in \mathbb{R}^n$, and $b \in \mathbb{R}$ is the bias term. More specifically, a perceptron receives an input vector $\mathbf{x} = (x_1, x_2, \ldots, x_n)^T$, where each component x_i is multiplied by a weight w_i. The sum of these weighted inputs, along with a bias term b constitutes the net input of the perceptron, which also called pre-activation value and noted as a. This net input a is then processed by a activation function (in the case of perceptron, activation function is usually a sign function) to decide the perceptron output z, which determines the class label. The sign function is defined as:

$$\mathtt{sign}(a) = \begin{cases} 1 & \text{if } a \geq 0 \\ -1 & \text{if } a < 0 \end{cases} \tag{5.2}$$

This function outputs 1, signifying the positive class, if the net input a is greater than or equal to zero; otherwise, it yields -1, signifying the negative class. Each input vector $\mathbf{x}$ is associated with a label y, which takes on values of either $+1$ or -1, indicating the class to which the input belongs. The prediction of perceptron z is then compared to the actual label y to determine whether the classification is correct. If the prediction z matches y, the perceptron has successfully classified the input. if not, when $y \cdot z \leq 0$, an error has occurred, the weights $\mathbf{w}$ and bias b are adjusted according to minimize the discrepancy. The learning process of perceptron is governed by a set of rules outlined in Algorithm 1.

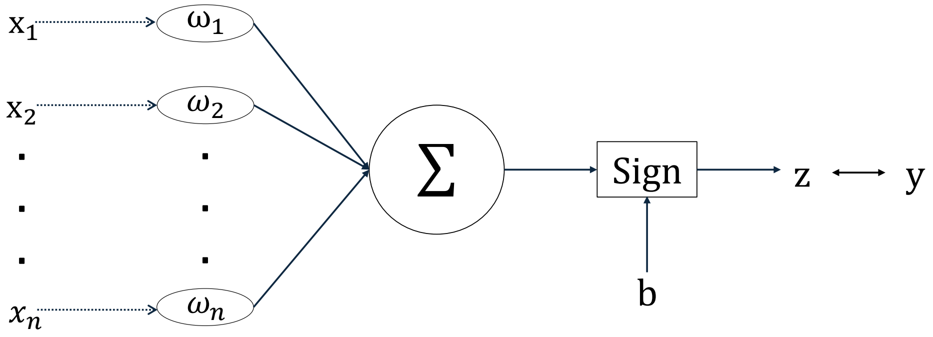

Figure 5.1: Diagram representation of the perceptron, illustrating its structure and the binary categorization process.

Algorithm 1 Perceptron Learning Rule

```
Initialize weights $\mathbf{w}$ and bias $b$ to zero or small random values.
for all training examples do
    Present input $\mathbf{x}$ and its label $y$ to the perceptron.
    Compute the linear combination: $a = \mathbf{w} \cdot \mathbf{x} + b$.
    if the perceptron makes an error, i.e., $y \cdot z \leq 0$ then
        Update the weights and bias: $\mathbf{w} \leftarrow \mathbf{w} + y \cdot \mathbf{x}$ and $b \leftarrow b + y$.
    end if
end for
```

5.2.2 Multi-layer perceptrons

Multi-layer perceptrons (MLPs) are an extension of the basic perceptron model, designed to overcome the limitation of linear separability by introducing one or more hidden layers (Werbos, 1974; Rumelhart et al., 1986). These hidden layers enable MLPs to model complex relationships in the data through non-linear transformations. The MLP architecture is comprised of an input layer, several hidden layers, and an output layer, with each layer consisting of multiple neurons. The neurons within each layer are interconnected to the subsequent layer through a network of weighted connections. Crucially, the introduction of an activation function in each neuron imparts non-linear characteristics to the network. The Figure 5.2 provide a visual representation of the MLP structure and its components.

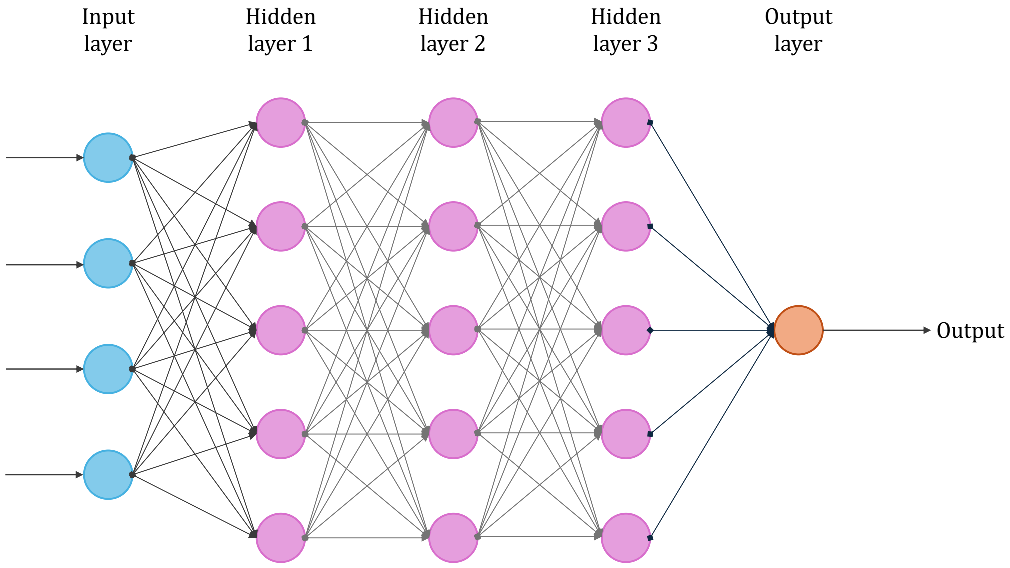

Figure 5.2: Diagram representing the flow of data through a multi-layer perceptron, comprising the input layer, hidden layers and output layer. The strength of each connection between neurons is defined by a weight, and each neuron incorporates a bias term. The activation function is applied element-wise in order to introduce non-linearity.

The mathematical operations within a MLP can be described as a series of transformations across layers of neurons. Each neuron in layer l computes its output based on the inputs it receives from the previous layer, modified by weights, an activation function, and a bias term. For any given neuron within layer l, the computation is represented by the following equation:

$$\mathbf{z}_l = f_l(\mathbf{z}_{l-1}) = \sigma_l(\mathbf{W}_l \cdot \mathbf{z}_{l-1} + \mathbf{b}_l) \tag{5.3}$$

This equation represents the process where the $\mathbf{z}_{l-1}$ from the preceding layer are transformed. $\mathbf{W}_l$ represents the matrix of weights connecting layer $l-1$ and layer l. Each element $w_{l,ji}$ within this matrix pertains to the weight of the connection from neuron i in layer $l-1$ to neuron j in layer l. $\mathbf{b}_l$ is the bias vector, where each element $b_{l,j}$ is the bias for neuron j in the layer l. $\sigma_l(\cdot)$ denotes the activation function, which is applied element-wise to the neuron's input. In addition, the input to the activation function is referred to as pre-activation value and can be noted with $\mathbf{a}_l = \mathbf{W}_l \cdot \mathbf{z}_{l-1} + \mathbf{b}_l$. Figure 5.3 illustrates a single layer within a MLP, highlighting the computation process in a layer with two neurons. The figure depicts the weighted inputs, bias terms, and activation functions that contribute to the output of each neuron. The arrows indicate the flow of information through the layer, where the inputs are transformed into outputs via the application of weights $\mathbf{W}$ and biases $\mathbf{b}$. The activation functions σ are applied to the pre-activation values $\mathbf{a}$ to generate the final outputs $\mathbf{z}$ for each neuron.

To focus on a single neuron j within layer l, the computation can be expressed in scalar form as:

$$z_{l,j} = \sigma_l\left(\mathbf{w}_{l,j}^T \mathbf{z}_{l-1} + b_{l,j}\right) \tag{5.4}$$

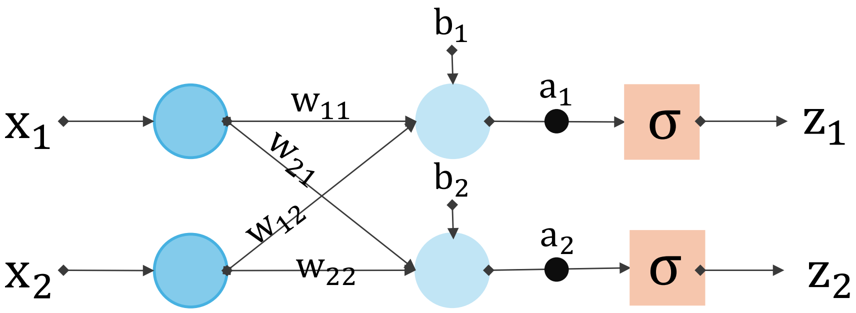

Figure 5.3: Diagram representing the computational process of a single MLP layer with two neurons, showing the flow of inputs through weights, biases, and activation functions to generate outputs.

This illustrates the sum of the weighted inputs from all neurons i in the previous layer, summed with the bias $b_{l,j}$, and then passed through the activation function.

The training process of a MLP relies on a algorithm known as back propagation, which is used to adjust the weights and biases in a neural network by minimizing the error between the predicted output and the actual labels, effectively guiding the learning process of the network. Specifically, an MLP consists of multiple layers of neurons, with the output of each layer being passed to the next layer through weighted connections. During training, after the input data has been propagated forward through the network, the resulting output is compared with the target labels, generating an error. The back-propagation algorithm calculates the gradient of this error with respect to each parameter using the chain rule, and adjusts the weights and biases accordingly, gradually improving the accuracy of the model's predictions. The details of the back propagation algorithm will be discussed in the subsequent Section 5.2.5. When we reach the final layer N, the formulation for the output layer is:

$$\mathbf{z}_N = f_N(\mathbf{z}_{N-1}) = \sigma_N(\mathbf{W}_N \cdot \mathbf{z}_{N-1} + \mathbf{b}_N) \tag{5.5}$$

Here, the pre-activation value for the output layer is: $\mathbf{a}_N = \mathbf{W}_N \cdot \mathbf{z}_{N-1} + \mathbf{b}_N$.

5.2.3 Activation function

If a MLP were simply a stack of perceptrons, each with an activation function that is not differentiable, such as the sign function or the Heaviside step function, this would make the entire model difficult to train, as the parameter updating relies on gradient calculations which are not available with non-differentiable functions. To address this issue, these non-differentiable functions can be replaced with non-linear, differentiable functions. In practical application, on such function that is widely utilized due to its differentiable nature is the sigmoid function (Rumelhart et al., 1986) defined as:

$$\sigma(x) = \frac{1}{1 + \exp^{-x}} \tag{5.6}$$

This function maps the input values to a range between 0 and 1, making it particularly useful for binary classification tasks where the output needs to represent a probability. Its derivative, $\sigma'(x) = \sigma(x) \cdot (1 - \sigma(x))$, is essential for the back propagation process, allowing for the efficient computation of gradients and the subsequent update of network weights.

Another widely used activation function in neural networks, especially in deep learning architectures, is the Rectified Linear Unit (ReLU) (Fukushima, 1980). The ReLU function is defined as:

$$\text{ReLU}(x) = \max(0, x). \tag{5.7}$$

This function introduces non-linearity into the model while being computationally efficient. It outputs the input value unchanged if it is positive; otherwise, it outputs zero. The simplicity of the ReLU function

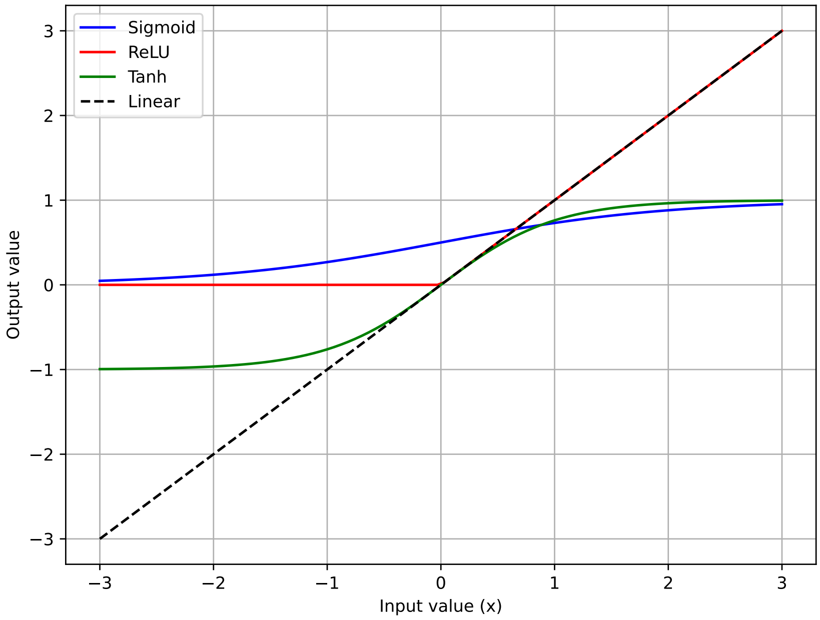

Figure 5.4: An illustration of commonly used activation functions.

contributes to its effectiveness in allowing neurons to fire more readily than with the Sigmoid or tanh functions, which can be beneficial for the training of deep networks. The derivative of the ReLU function is straightforward:

$$\mathrm{ReLU}'(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{5.8}$$

This piece-wise function indicates that during back propagation, if the neuron output is positive, the gradient is passed through unaltered; if not, the gradient is zeroed.

Figure 5.4 illustrates the behavior of three distinct activation functions: the sigmoid, the ReLU, and the tanh function. The sigmoid function exhibits saturation at 1 for positive inputs and 0 for negative inputs, which can lead to the vanishing gradient problem during training. The tanh function is similar but centered at -1 and +1, yet also susceptible to saturation. In contrast, the ReLU function remains non-saturating for positive inputs, meaning its gradient is constant and does not vanish. This characteristic facilitates more stable gradient propagation in deep networks. However, it can encounter the "dying ReLU" issue, where neurons produce zero outputs for all inputs due to negative weights, effectively deactivating them. The linear function $y = x$ is a straightforward mapping without introducing non-linearity, making it less suitable for complex models. The choice between these functions is critical for the model ability to learn and generalize from data.

5.2.4 Loss function

The loss function plays a vital role within the MLP, evaluating the deviation between the predictions of the model and the actual target values. During the training phase, the loss function calculates a loss value for each training sample, with the overall loss for the entire dataset being obtained by summing or averaging these individual values. The aggregated loss function provides a quantitative metric for evaluating the performance of the model.

Beyond performance evaluation, the loss function is also a key driver in the gradient descent optimization process. By calculating the gradients of the loss function with respect to the model parameters, we can determine adjustments to the weights and biases that minimize prediction errors and enhance the accuracy of model.

5.2.5 Back propagation

Back propagation is a key algorithm in training MLPs, which calculates the gradient of the loss function with respect to the network parameters, enabling the update of weights and biases (Werbos, 1974; Rumelhart et al., 1986). The process involves two main phases: forward propagation and back propagation.

Given an MLP comprising N layers, the input data to the network is x. The function formulation at the n-th layer, f_l, defines how the inputs are transformed within that layer.

Forward propagation is the forward process of data flow in a neural network. During this process, the input data first enters the first layer of the network, and then passes through each layer with weighted sums and possible non-linear activation function transformations until the last layer, generating the network output. This process can be seen as a series of matrix multiplications and function applications, aimed at extracting features from the input data for prediction or classification. Forward propagation is the first step in the training and inference of a neural network, determining the network response to input data.

$$\hat{y} = f_N \circ f_{N-1} \circ \dots f_1(x) \tag{5.9}$$

Once the forward propagation is complete and the output $\hat{y}$ is obtained, the next step is to calculate the error of the network's predictions with respect to the true values. This is typically done using a loss function L, which measures the discrepancy between the predicted outputs $\hat{y}$ and the actual target values y. The loss L is computed as : $L(y, \hat{y})$.

Following this, the backward pass, also known as back propagation, involves computing the gradient of the loss function with respect to each parameter in the network, starting from the output layer and moving backward to the input layer. The backpropagation algorithm is instrumental in calculating these gradients. Armed with these gradients, the gradient descent method is then employed to update the parameters. Gradient descent is an optimization technique that iteratively moves toward the minimum of a function, in this case, the loss function L. The update rule for each parameter, such as weights $\mathbf{W}$ and biases $\mathbf{b}$ is given by:

$$\mathbf{W}' \leftarrow \mathbf{W} - \eta \cdot \nabla_{\mathbf{W}} L \tag{5.10}$$

$$\mathbf{b}' \leftarrow \mathbf{b} - \eta \cdot \nabla_{\mathbf{b}} L \tag{5.11}$$

Here, η is the learning rate, a hyper parameter that determines the step size at each iteration. By adjusting the parameters in the direction that reduces the loss, gradient Descent aims to find the set of parameters that minimizes the loss function, thus optimizing the model's predictive performance.

This is done using the chain rule of calculus, which allows the computation of the gradient of the loss with respect to the inputs and parameters of each layer. For the output layer N, according to the Equation(5.5), the gradient of the loss function L with respect to the post-activation values $\mathbf{z}_N$, denoted as $\frac{\partial L}{\partial \mathbf{z}_N}$, can be directly computed with the loss function being used in the layer. Then we can calculate the gradient of the pre-activation values a_N for the output layer. This requires the chain rule and the derivative of the activation function f_N and can be formulated as:

$$\frac{\partial L}{\partial \mathbf{a}_N} = \frac{\partial L}{\partial \mathbf{z}_N} \cdot \sigma'_N(\mathbf{a}_N) \tag{5.12}$$

The gradient for weights $\mathbf{W}_N$ and biases $\mathbf{b}_N$ can be computed as:

$$\frac{\partial L}{\partial \mathbf{W}_N} = \frac{\partial L}{\partial \mathbf{a}_N} \cdot \frac{\partial \mathbf{a}_N}{\partial \mathbf{W}_N} \tag{5.13}$$

$$= \left(\frac{\partial L}{\partial \mathbf{z}_N} \odot \sigma'_N(\mathbf{a}_N) \right) \cdot \mathbf{z}_{N-1}^T \tag{5.14}$$

In this equation, the symbol $\odot$ denotes as the Hadamard product, which signifies that each element of the gradient of the loss with respect to $\mathbf{z}_N$ is multiplied by the corresponding element of the derivative of the activation function $\sigma'_N(a_N)$. The result is then multiplied by the transpose of the activations from the previous layer $\mathbf{z}_{N-1}^T$. For the biases $\mathbf{b}_N$, the expression is given by:

$$\frac{\partial L}{\partial \mathbf{b}_N} = \frac{\partial L}{\partial \mathbf{a}_N} \cdot \frac{\partial \mathbf{a}_N}{\partial \mathbf{b}_N} \tag{5.15}$$

$$= \frac{\partial L}{\partial \mathbf{z}_N} \odot \sigma'_N(\mathbf{a}_N) \tag{5.16}$$

For the hidden layer $l \leq N$, the process is recursive. Assume that we have already computed the gradients $\frac{\partial L}{\partial \mathbf{z}_{l+1}}$ for the layer $l+1$. We use these gradients to compute the gradients for layer l.

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{z}_l} &= \frac{\partial L}{\partial \mathbf{z}_{l+1}} \cdot \frac{\partial \mathbf{z}_{l+1}}{\partial \mathbf{z}_l} & (5.17) \\ &= \frac{\partial L}{\partial \mathbf{z}_{l+1}} \cdot \mathbf{W}_{l+1} \odot \sigma'_{l+1}(\mathbf{a}_{l+1}) & (5.18) \end{aligned}$$

and then, for layer l, the gradient for weights $\mathbf{W}_l$ and biases $\mathbf{b}_l$ are articulated as:

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{W}_l} &= \left(\frac{\partial L}{\partial \mathbf{z}_l} \odot \sigma'_l(\mathbf{a}_l) \right) \cdot \mathbf{z}_{l-1}^T & (5.19) \\ \frac{\partial L}{\partial \mathbf{b}_l} &= \frac{\partial L}{\partial \mathbf{z}_l} \odot \sigma'_N(\mathbf{a}_l) & (5.20) \end{aligned}$$

5.2.6 Optimization

Following the exposition of the back propagation algorithm in the preceding section, this section examines the various methods available for optimizing the parameters of MLPs. Although the gradient descent method, as previously discussed, is the foundation for such optimization, it is not without its limitations. The objective of this section is to examine the shortcomings of gradient descent and to present alternative optimisation techniques that have been developed to address these limitations.

Gradient descent represents a fundamental optimisation method, whereby the parameters of a multilayer perceptron are updated in a direction opposite to that of the gradient of the loss function. Nevertheless, its fundamental form is subject to certain inherent constraints. The computational cost of calculating gradients across the entire dataset can be considerable, particularly with large datasets, which can result in a slow optimisation process. Furthermore, the necessity to consider the entire dataset when making updates can result in delayed adjustments to the learning rate, which may further impede convergence. Furthermore, gradient descent is susceptible to becoming trapped in local minima or saddle points, which can considerably extend the training process. These challenges have prompted the development of alternative optimisation techniques with the aim of enhancing efficiency, accelerating convergence and providing a more effective means of navigating the complex loss surfaces encountered in training neural networks.

In order to tackle these shortcomings, the stochastic gradient descent (SGD) algorithm (Robbins and Monro, 1951) was introduced. In contrast to the conventional gradient descent approach, the SGD method updates the parameters based on a single training example (or a small batch) at a time, resulting in accelerated computation and reduced memory usage. The algorithm introduces controlled noise into the learning process, which can facilitate the escape from local minima and saddle points. To simplify the subsequent discussion, we will use θ to represent the collective set of parameters that are optimized during training, involving both the weight $\mathbf{W}$ and biases $\mathbf{b}$.

$$\theta \leftarrow \theta - \eta \cdot \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \nabla_\theta L(\theta; x^{(i)}, y^{(i)}) \tag{5.21}$$

where $|\mathcal{B}|$ is the number of samples in the mini-batch. The learning rate, denoted by η , is a crucial hyper parameter that dictates the magnitude of each update step in the optimization process. The gradient $\nabla_\theta L(\theta)$ represents the vector of partial derivatives of the loss function L with respect to all the parameters θ.

The Momentum method (Nesterov, 1983) is an optimization technique that enhances the SGD algorithm by incorporating a momentum term into the parameter updates. This technique helps to expedite the convergence of the algorithm by mitigating the fluctuations that occur during the gradient descent process. The momentum term calculates a velocity vector that accumulates the impact of past gradients, allowing the updates to reflect not only the current gradient but also the direction of previous updates. The mathematical expression for the Momentum method is represented as follows:

$$\begin{aligned} v &\leftarrow \gamma v + \eta \cdot \nabla_\theta L(\theta) & (5.22) \\ \theta &\leftarrow \theta - v & (5.23) \end{aligned}$$

In this expression: v denotes the velocity vector, which aggregates the gradients over time, γ $(0 < \gamma \leq 1)$, is the momentum coefficient, indicating the proportion of the previous velocity to be preserved. η represents the learning rate, a hyperparameter that scales the update step, and $\nabla_\theta L(\theta)$ signifies the gradient of the loss function L with respect to the parameters θ.

RMSProp is an adaptive learning rate optimization algorithm, proposed by Hinton (2012). It is designed to overcome the weakness of the standard SGD and is particularly useful for non-stationary objectives and for neural networks with long training times. RMSProp adjusts the learning rate for each parameter by dividing the gradient by an exponentially decaying average of squared gradients. The update rule for RMSProp is:

$$\theta \leftarrow \theta - \frac{\eta}{\sqrt{\epsilon + \rho \cdot g_{\text{prev}}^2}} \cdot \nabla_\theta L(\theta) \tag{5.24}$$

where ϵ is a small constant to prevent division by zero, ρ is the decay rate for the squared gradients. g_{prev}^2 is the exponentially weighted average of squared gradients from the previous updates.

Adam (Kingma, 2014), which stands for Adaptive Moment Estimation, is another adaptive learning rate method that combines the advantages of Momentum method and RMSProp algorithm. It computes the exponentially weighted average of both the gradient (for momentum) and the squared gradient (for RMSProp), providing a method that can adaptively adjust the learning rate for each parameter. The mathematical expression of Adam is given by:

$$m_t \leftarrow \beta_1 m_{t-1} + (1 - \beta_1) \nabla_\theta L(\theta), \tag{5.25}$$

$$v_t \leftarrow \beta_2 v_{t-1} + (1 - \beta_2)(\nabla_\theta L(\theta))^2, \tag{5.26}$$

$$\theta_{t+1} \leftarrow \theta_t - \frac{\eta}{\sqrt{\frac{v_t}{1-\beta_2^t}} + \epsilon} \cdot \frac{m_t}{1 - \beta_1^t}, \tag{5.27}$$

with m is the exponentially weighted average of the gradients, v is the exponentially weighted average of the squared gradients. β_1 and β_2 are the decay rates for the first and second moment estimates, respectively.

Both RMSProp and Adam are widely used in deep learning due to their ability to adapt the learning rate based on the parameters being updated, which often results in faster convergence and better performance than traditional SGD. For the training of the proposed networks, which will be discussed in Section 7.1.1, we employed the Adam algorithm for optimization. The primary reasons for this choice are that Adam combines the advantages of the Momentum method and the RMSprop algorithm, which can accelerate gradient descent and reduce oscillations. Additionally, its adaptive learning rate mechanism enables faster convergence during training, saving time and enhancing efficiency.

5.3 Convolutional neural networks

Convolutional neural networks (CNN) represent a revolutionary technology within the field of deep learning, achieving remarkable success in image processing and computer vision tasks. The origins of the CNN can be traced back to the 1960s, when the first simple neural network models were developed with the aim of mimicking visual perception. However, the resurgence of neural networks and the advent of back propagation algorithms in the 1980s marked a turning point, as this was when convolutional networks began to receive greater research and development attention. The concept of CNN came from LeNet proposed by LeCun et al. (1998), which was successfully applied to handwritten digit recognition. The significance of LeNet lies in its introduction of convolutional and pooling layers, recognised as fundamental concepts in the design of convolutional neural networks. Figure 5.5 depicted the architecture of LeNet. This work marked the early success of convolutional neural networks and set the stage for future research. Over time, especially after the victory of AlexNet in the 2012 ImageNet competition, CNN began to emerge as the tool of choice for image recognition and other visual tasks.

A defining characteristic of CNN is their ability to automatically and effectively identify and learn features from images, such as edges and textures, without the need for manual feature engineering. This marks a significant advancement over earlier image processing methods. Key features of CNN include convolutional layers that apply learnable filters to the input image, activation functions such as ReLU that introduce non-linearity, pooling layers that reduce the spatial dimension of features and enhance

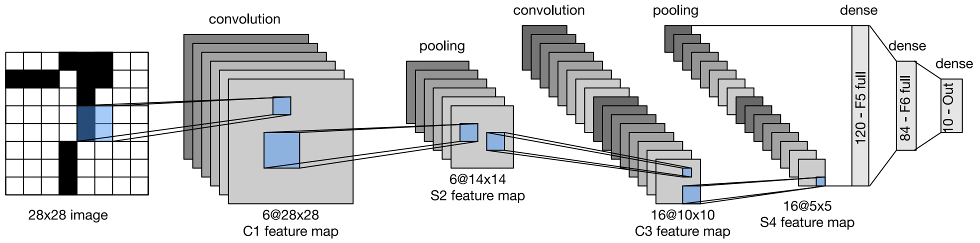

Figure 5.5: Network Architecture and Data Flow of LeNet. The input is an image of handwritten digit, the output layer delivers probabilities for the 10 potential digits. The network's architecture is composed of 2 convolutional layers, 2 pooling layers and a fully connected layer. Figure is extracted from Zhang et al. (2023).

invariance, and fully connected layers that integrate these features for final classification or regression tasks.

CNN are essentially a specialized form of MLPs, yet their architecture capitalizes on the two-dimensional structure of images through convolutional and pooling layers, enabling parameter sharing and reducing computational complexity. This design renders CNN particularly well-suited for processing data with grid-like topology, such as images, videos or spectrograms.

A typical convolutional neural network (CNN) comprises several main components that work in concert to form a powerful model capable of addressing complex visual problems across various domains. The network begins with the input layer, which receives the raw pixel values of an image. It then proceeds through alternating convolutional and activation layers, which are designed for feature extraction and transformation. Subsequently, pooling layers are employed to reduce the dimensionality of the features, thereby distilling the most salient characteristics. It is possible to include normalization layers, such as batch normalization, if desired, with the aim of accelerating training and enhancing model stability. At the conclusion of the network, fully connected layers integrate the extracted features for final classification or other tasks. This process completes in the output layer, which generates predictions based on the network's objective, whether classification, regression, or other forms of analysis.

5.3.1 Convolutional layers

Following the historical context of CNN, we now turn our attention to one of their core components—the convolutional layer. The convolutional layer is the building block of CNN, responsible for the automatic feature extraction that makes these networks so powerful in processing grid-like data structures such as images. For a visual representation of the convolution process, Figure 5.6 illustrates the convolution process in detail, where a 3×3 input matrix undergoes a convolution with 2×2 filter to produce a 2×2 output. This simplified example highlights how local features are extracted through filter application.

At the heart of a CNN is the convolutional operation, which involves the use of a kernel or filter that slides over the input image to produce a feature map. This operation constitutes a mathematical process whereby a dot product is applied between the kernel and the local region of the image it covers. This results in the highlighting of specific features, such as edges, textures, or more complex patterns, depending on the depth of the network. The operation of a convolutional layer involves a series of steps centered around the application of a kernel, a compact matrix of weights, to the input data. As the kernel moves across the input in a process known as convolution, it computes the sum of products between the kernel elements and the local regions of the input, revealing features such as edges or textures. The use of multiple kernels in parallel allows for the simultaneous detection of various features, with each kernel generating a distinct feature map. The total depth of the output is determined by the number of kernels, or filters, used.

Key parameters that dictate the behavior of convolutional layers include the kernel size, typically a small square like 3×3 or 5×5; the stride, which defines the kernel's movement across the input; and padding, which adds zeros around the input to adjust the spatial dimensions of the output feature map. The number of filters corresponds to the number of kernels and dictates how many feature maps are produced.

The output size of the feature map is determined by the input size and several operational parameters:

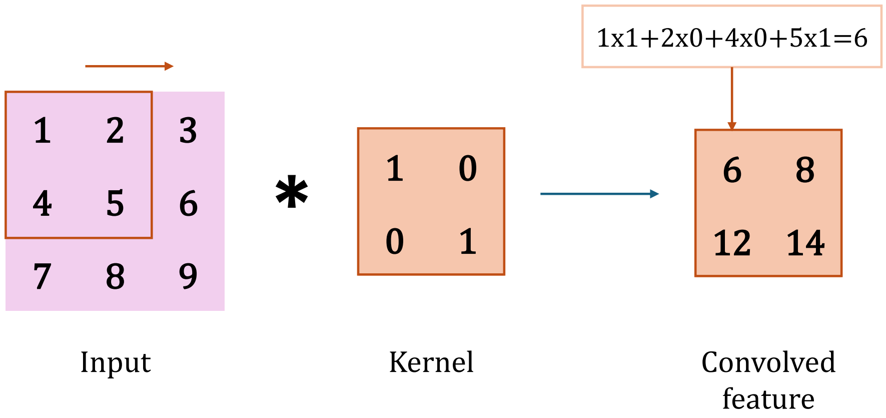

Figure 5.6: Demonstration of the convolution operation in a CNN. A 3×3 input matrix is convolved with a 2×2 kernel using a stride of 1, resulting in a 2×2 output.

the kernel size, padding, and stride. This relationship is summarized in the following formula:

$$\text{Output size} = \left\lfloor \frac{\text{Input size} + 2 \times \text{Padding} - \text{Kernel size}}{\text{Stride}} + 1 \right\rfloor \tag{5.28}$$

This equation ensures that the dimensions of the output feature map are calculated accurately, reflecting the combined effects of the convolutional parameters.

5.3.2 Pooling layer

After the convolutional layers have extracted the initial features from the input data, pooling layers in CNN seamlessly advance to perform a critical downsampling operation. This systematic reduction in the spatial dimensions of the feature maps serves a dual purpose: it alleviates the computational burden and instills translation invariance, an essential property that enables the network to recognize features irrespective of their position within the input space.

The max pooling operation, widely utilized in CNN, operates by defining a pooling window, typically a 2x2 or 3x3 square, which moves across the feature map with a specified stride. As the window traverses the map, the max pooling function identifies and retains the maximum value within each region, effectively capturing the most prominent feature. This process ensures that the most significant information is preserved while redundant details are discarded, thereby enhancing robustness of the network and reducing its sensitivity to variations in scale and orientation.

By producing a down sampled feature map with reduced dimensions, pooling layers significantly decrease the computational load for subsequent layers. Moreover, they confer a degree of feature invariance, allowing the network to maintain its recognition capabilities despite variations in the scale and position of features within the input. Figure 5.7 provides a visual representation of the operation of a max pooling layer, demonstrating how it selects the maximum value within each defined window, thereby resulting in a feature map with reduced spatial dimensions.

5.3.3 Applications of CNN in vision tasks

CNN has profoundly transformed the field of image processing, providing unprecedented capabilities in analyzing and interpreting visual data, and demonstrating exceptional performance across various visual tasks. such as image classification, object detection, and image segmentation. Inspired by the human visual system, these networks leverage convolutional layers and pooling layers to efficiently and adaptively learn spatial hierarchies of features from input images. As mentioned earlier in this sections, LeNet has been proposed for handwritten digit recognition in the 1990s. LeNet did not receive much attention in the following years due to the lack of large-scale datasets and computational resource constraints. However,

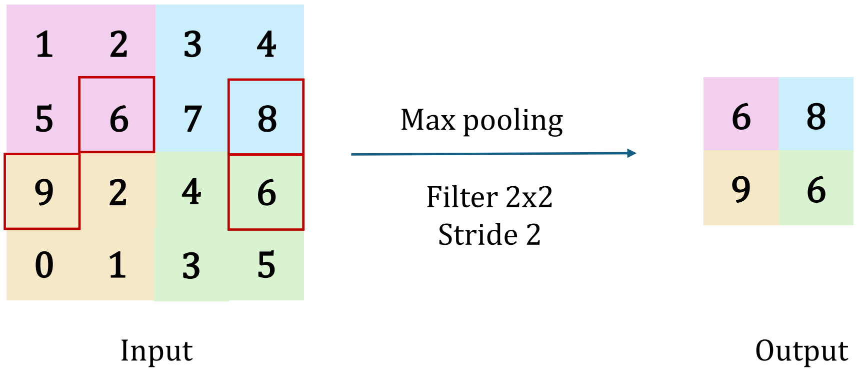

Figure 5.7: Demonstration of the max pooling operation in a CNN. A 2×2 pooling window with a stride of 2 (represented by different color in the input) is applied to a 4×4 input feature map, producing a downsampling 2×2 output feature map by retaining the maximum values (red box) from each pooling window (indicated by different colored backgrounds).

with the advent of ImageNet (Deng et al., 2009) and the efficient computational power of GPUs, the conditions that previously limited LeNet impact have been addressed. In 2012, AlexNet (Krizhevsky et al., 2012) made a significant impact on the field of computer vision by introducing the ReLU activation function and the Dropout technique. It also fully benefited the parallel computing capabilities of GPUs, which accelerated the training process. AlexNet achieved breakthrough results in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC), becoming the first deep neural network to outperform traditional computer vision methods. This victory marked the beginning of a new era for convolutional neural networks in visual tasks.

Image classification

Image classification is a fundamental task in computer vision that involves assigning input images to specific categories or labels. This task can be thought of as a pattern recognition problem, where a computer learns to recognise and classify different objects or scenes based on the visual features of an image. For example, given an image of a cat, the model should correctly classify it as "cat" among various possible categories like dog, cat, or bird. Figure 5.8 illustrates a simplified workflow of image classification using a CNN. The input is an image of a cat, which is processed through a single convolutional layer followed by a pooling layer. The result is then passed to a fully connected layer that outputs the probabilities of the image being classified into different animal categories (e.g., cat, dog, bird, and tiger). In mathematical terms, the process of image classification is formulated to identify a function that can map an input image x of dimensions $m \times n \times c$, where m and n are the height and width of the image, and c is the number of channels(e.g., 1 for gray images, 3 for RGB images) to its corresponding class label over k possible classes with a certain level of confidence. The mathematical expression for this process is represented as shown, and Figure 5.8 illustrates the overall process of image classification, from input image to predicted class label.

$$\hat{\mathbf{y}} = f(\mathbf{X}, \theta) \tag{5.29}$$

where $\mathbf{X}$ represents the input image, f the operation of the neural network. $\hat{\mathbf{y}}$ is the predicted probability vector, $\theta = \{W, b\}$ represents the learnable parameters of the model.

After the success of AlexNet in the field of image classification, numerous neural network architectures have been developed to further enhance the performance and capabilities of deep learning models in this domain. VGG (Simonyan and Zisserman, 2014) build deep neural networks by stacking multiple small convolutional kernels (3×3), reducing the number of parameters while maintaining the model's power. ResNet (He et al., 2016) introduced residual connections design, enabling the training of extremely deep networks and significantly improving the performance and generalization of the models. Google's Inception module innovatively applied different sized convolutional kernels (1×1,3×3, 5×5) and pooling operations in parallel on the same layer, enhancing the network's ability to extract features

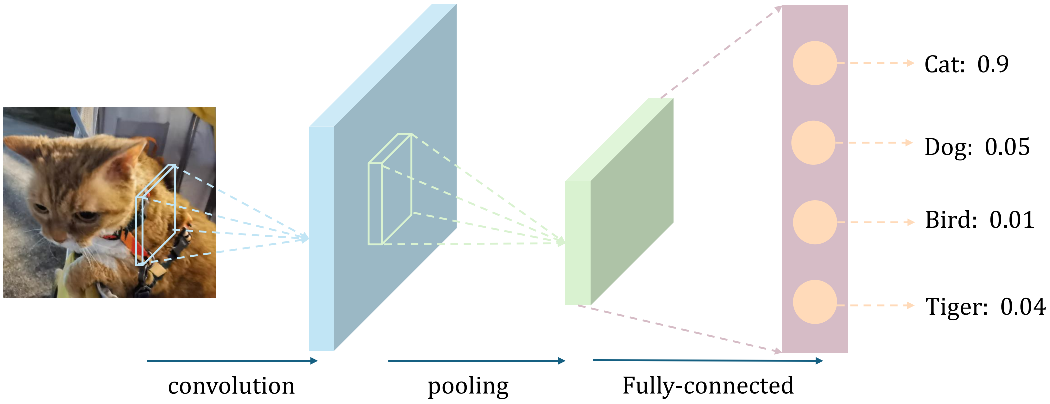

Figure 5.8: Simplified illustration of a CNN workflow for image classification. The input cat image is processed through a convolutional layer, followed by a pooling layer, and finally classified into different animal categories based on output probabilities. (This diagram is a conceptual representation for illustrative purposes, highlighting only the basic components of the CNN.)

at different scales (Szegedy et al., 2015). Since then, the Inception module has been improved with Batch Normalization techniques (Ioffe, 2015), more efficient Inception module (Szegedy et al., 2016), and residual block (Szegedy et al., 2017). DenseNet (Huang et al., 2017) connects each layer to all previous layers to increase the efficiency of information exchange and parameter utilization in the network. Based on these advancements, EfficientNet (Tan, 2019) employs a systematic approach to scaling the depth and width of the network, optimizing model performance while maintaining computational efficiency through a compound scaling method. These advancements in image classification networks have not only furthered the field of image classification itself, but have also significantly accelerated progress in many other tasks, such as object detection and image segmentation. The main structure of these networks, with their powerful feature extraction capabilities, can be effectively utilized as a backbone for a diverse array of tasks. Inspired by these advancements, this work incorporates several of these influential design elements into the proposed network architecture. For instance, we have adopted small convolutional kernels, ReLU activation functions, and residual structure to enhance the network performance. The detailed design of network, as well as the implementation and impact of residual layers, will be thoroughly introduced in subsequent sections. The detailed design of the network as well as the specific introduction of the residual layer will be presented in the later Section 6.3.

Image segmentation

Image segmentation is a core task in computer vision, where the goal is to divide an image into meaningful segments or regions, typically corresponding to different objects or areas of interest within the image. Unlike image classification, which assigns a single label to the entire image, image segmentation involves labelling each pixel in the image, providing a more detailed understanding of the image. There are two main types of image segmentation: semantic segmentation, where each pixel is assigned a label based on the object class it belongs to, and instance segmentation, where different instances of the same object class are distinguished and labelled separately.

While image classification assigns a single label to an entire image, image segmentation is a more granular task that involves assigning a label to each individual pixel in the image. In other words, instead of mapping an image x to a single class label, image segmentation maps each pixel (i, j) in the image to class label $y_{i,j}$. Thus the image segmentation problem can be represented as :

$$\hat{\mathbf{Y}} = f(\mathbf{X}, \theta) \tag{5.30}$$

where $\hat{\mathbf{Y}}$ is the output segmentation map, with the same spatial dimensions as the input image $\mathbf{X}$. Each element $\hat{y}_{i,j}$ in $\hat{\mathbf{Y}}$ corresponds to the predicted class label for the pixel at position (i, j) in $\mathbf{X}$. The concept of image segmentation is depicted in Figure 5.9

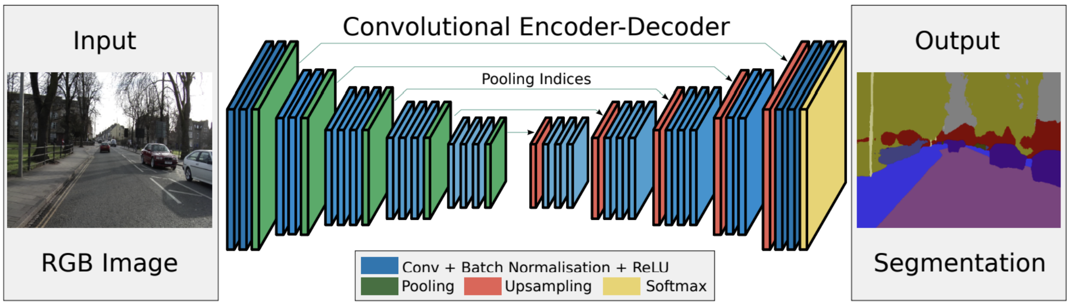

Figure 5.9: Illustration of a CNN with an encoder-decoder structure workflow for image segmentation. The RGB input image is subjected to downsampling to condense its spatial resolution, followed by upsampling to restore it to the original scale. The resultant output is a detailed feature map that corresponds to the segmented components of the input image. Figure extracted from Badrinarayanan et al. (2017)

In 2014, FCN (Long et al., 2015) proposed an end-to-end network architecture for the first time, extending convolutional neural networks from traditional image classification to pixel-level image segmentation. U-Net (Ronneberger et al., 2015) adopted a unique U-shaped encoder-decoder network structure, which contains a downsampling path (encoder) and a symmetric upsampling path (decoder). This network architecture design allows the network to extract information during downsampling and recover image details during upsampling, while connecting the information in the upsampling and downsampling paths through hopping connections, allowing more positional information and details to be retained during the segmentation process. Although U-Net was initially designed specifically to solve the medical image segmentation problem, its excellent performance on the segmentation problem has led to its rapid and widespread use in a variety of image segmentation tasks. SegNet Badrinarayanan et al. (2017) employs a similar encoder-decoder architecture, but instead of direct connections between encoder and decoder layers, it uses pooling indices during upsampling, reducing computational overhead while maintaining segmentation accuracy. DeconvNet (Noh et al., 2015) employs transposed convolutions to reconstruct the spatial resolution during the upsampling satge, which helps to produce more detailed segmentation maps, though at the cost of increased computational complexity. The DeepLab series progressively advanced semantic segmentation by introducing atrous convolutions (Chen et al., 2014) and Atrous Spatial Pyramid Pooling (ASPP) (Chen et al., 2017) for multi-scale context, while later versions, like DeepLabV3 (Chen, 2017) and DeepLabv3+ (Chen et al., 2018), adopted an encoder-decoder structure for better boundary precision, eliminating the need for post-processing steps like CRFs. These advancements in the segmentation field also provide valuable insight into other image-to-image tasks, such as image generation (Isola et al., 2017) and image restoration. In the field of image restoration, techniques such as the encoder-decoder architecture, U-shaped structure, and deconvolution layer have been proven to be highly effective for tasks such as image denoising and super-resolution. These techniques have also been integrated into the proposed network architecture, which will presented later in the Section 6.3 with the aim of enhancing the performance of model on image restoration tasks.

Image restoration

Image restoration is another important challenge in the field of computer vision, which focuses on recovering degraded or damaged images. Degradation can arise from various factors, including noise, blur, or missing data, which can significantly impact the quality and usability of the images. The primary goal of image restoration is to reverse these effects and reconstruct a clean image that closely resembles the original, undisturbed version. This process is often framed as an inverse problem, where the objective is to estimate the original image $\mathbf{X}$ from a degraded observation $\mathbf{Y}$. The degradation process is typically modeled as:

$$\mathbf{Y} = \mathcal{D}(\mathbf{X}) + \mathbf{N} \tag{5.31}$$

where $\mathcal{D}(\cdot)$ represents the degradation operator, and $\mathbf{N}$ denotes the noise. An image restoration problem can be specified in more details by defining a particular degradation function $\mathcal{D}(\cdot)$, which represents the process of image corruption. For instance, in the context of image denoising, the degradation function $\mathcal{D}(\cdot)$ is usually omitted and the acquired image is modeled as the original clean image corrupted by an additive noise. Along the last years, denoising has been shown to be an essential building block to perform more advanced tasks (Milanfar and Delbracio, 2024). When considering image deblurring situation, the degradation function $\mathcal{D}(\cdot)$ describes the conversion of a clear image into a blurred image. This degradation is usually caused by the imaging system (e.g., due to the lens), the acquisition (e.g., camera shake, inaccurate focus) or the observed scene itself (e.g., fast moving objects). The goal of image deblurring is to recover a clear original image from an observed blurred image. In most works this degradation can be described as a spatially-invariant linear filtering and the function $\mathcal{D}(\cdot)$ stands for the convolution by a blurring kernel.

To solve image restoration problem problem, one generally seeks a function $f(\cdot)$ that estimates the clean image $\hat{\mathbf{X}}$ from the degraded image $\mathbf{Y}$, i.e., $\hat{\mathbf{X}} = f(\mathbf{Y}, \theta)$. When this function is a CNN, the network is parametrized by θ and it is typically trained by learning from examples of degraded images and corresponding clean images. This function is typically learned through a deep learning model, which is trained using pairs of degraded and corresponding clean images. In recent years, significant progress has been made in the image restoration field, with the application of deep neural network. For instance, Lefkimmiatis (2017) introduced N3Net, a CNN leveraging non-local operations for color image denoising, enhancing performance by capturing long-range dependencies within images. In particular, when considering image denoising, Zhang et al. (2017) proposed DnCNN that uses residual learning to outperform traditional Gaussian denoisers and set a new benchmark for this task. Zhang et al. (2018) presented FFDNet, a solution designed for fast and flexible image denoising through the use of grouped convolutions and depthwise separable convolutions. More recently, the residual dense network proposed by Zhang et al. (2020) further advanced the development of image restoration techniques with its innovative network structure. Finally, DRUNet was proposed by Zhang et al. (2021) and utilizes this deep denoiser to achieve exceptional results. It is worth mentioning here that the contribution reported in this manuscript elaborates on the architecture of DRUNet to perform RFI mitigationis a prototype of my proposed network architecture, and I have customized DRUNet so that it can be applied to our RFI mitigation problem (see Section 6.3).

Résumé du chapitre

L'intelligence artificielle a connu une révolution grâce à l'apprentissage profond, un domaine qui a transformé de nombreux secteurs en raison de sa capacité à extraire des motifs complexes à partir de grandes quantités de données. Les fondements de l'apprentissage profond reposent sur les réseaux de neurones artificiels qui, avec leurs architectures multi-couches, tentent d'imiter la complexité du cerveau humain. Cette structure permet aux réseaux de neurones profonds de se distinguer dans des tâches nécessitant une reconnaissance de motifs et une extraction de caractéristiques avancées. Le succès de l'apprentissage profond résulte de la combinaison de trois éléments clés : la disponibilité d'ensembles de données massifs pour l'entraînement des modèles, des progrès en calcul pour traiter ces données, et des avancées algorithmiques pour entraîner efficacement ces réseaux complexes. Ensemble, ces facteurs ont poussé l'apprentissage profond à de nouveaux sommets, dans des applications allant de la reconnaissance d'image et vocale au traitement du langage naturel.

Les réseaux de neurones, qui forment la base de l'apprentissage profond, sont composés de neurones interconnectés qui traitent les informations à travers des entrées pondérées et des fonctions d'activation. Chaque neurone reçoit des entrées provenant d'autres neurones, les traite à travers une fonction mathématique, puis transmet la sortie aux neurones suivants. Un neurone est caractérisé par un ensemble de poids qui détermine l'impact de chaque entrée sur sa sortie, ainsi qu'un biais qui ajuste le seuil d'activation.

Le perceptron (Rosenblatt, 1958), introduit dans les années 1950, est la forme la plus simple de réseau de neurones, agissant comme un classifieur linéaire pour deux classes. Il est inspiré des réseaux de neurones biologiques et fonctionne avec un vecteur d'entrée, des poids et un biais, utilisant une fonction d'activation signe pour déterminer la classe de l'entrée.

Les réseaux de neurones multicouches (MLP) (Werbos, 1974) généralisent le modèle du perceptron en ajoutant une ou plusieurs couches cachées, ce qui permet de modéliser des relations de données complexes via des transformations non linéaires. Ces couches cachées, combinées avec les fonctions d'activation, apportent des caractéristiques non linéaires au réseau.

Les fonctions d'activation sont essentielles dans les réseaux de neurones car elles introduisent la non-linéarité, permettant au réseau d'apprendre et d'effectuer des tâches complexes. Les fonctions d'activation les plus courantes sont la sigmoide et l'unité linéaire rectifiée (ReLU).

Les fonctions objectif sont cruciales pour évaluer la différence entre les prédictions d'un modèle et les valeurs réelles. Elles fournissent une mesure de la performance du modèle et guident l'optimisation en calculant les gradients par rapport aux paramètres du modèle.

La rétropropagation est un algorithme essentiel pour l'entraînement des réseaux de neurones multicouches (Rumelhart et al., 1986), car il calcule le gradient de la fonction objectif par rapport aux paramètres du réseau, permettant ainsi de mettre à jour les poids et les biais. Ce processus comprend deux phases : la propagation avant, où les données traversent le réseau pour produire une sortie, et la rétropropagation, où les gradients sont calculés pour ajuster les paramètres.

Les réseaux neuronaux convolutifs (CNN) ont été une révolution dans l'apprentissage profond, en particulier pour le traitement d'image et la vision par ordinateur. Ils ont été conçus pour identifier et apprendre automatiquement des caractéristiques d'images, tels que les contours et les textures, sans nécessiter d'ingénierie de caractéristiques manuelle. Les caractéristiques clés des CNN comprennent des couches convolutives avec des filtres ajustables, des fonctions d'activation comme la ReLU, des couches de mise en commun pour réduire la dimension et améliorer l'invariance, et des couches entièrement connectées pour la classification ou la régression.

Les applications des CNN en vision sont nombreuses, allant de la classification d'images à la segmentation et à la restauration d'images. Les CNN ont considérablement progressé dans ces domaines, offrant des outils puissants pour l'analyse des données visuelles.

Avec les avancées de l'apprentissage profond, les performances des CNN ont continué à progresser dans ces domaines, repousant les limites dans la résolution des tâches d'analyse d'image et de vision par ordinateur. Des architectures comme AlexNet, VGG, ResNet, Inception et DenseNet ont amélioré les performances et les capacités des modèles. Ces réseaux ont non seulement permis des progrès dans la classification d'images, mais ont également fourni des capacités d'extraction de caractéristiques puissantes pour d'autres tâches. Des architectures telles que U-Net et DeepLab ont fait des percées dans la segmentation d'images, tandis que des réseaux comme DRUNet ont montré des performances excellentes dans la restauration d'images. Ces développements offrent des outils puissants pour résoudre des problèmes pratiques et préparent le terrain pour de futures recherches et applications

Chapter 6

Deep learning for RFI mitigation

Convolutional neural networks (CNN) have demonstrated remarkable success in image processing tasks, establishing themselves as the preferred solution for a diverse array of computer vision applications. These networks, with their deep hierarchical structures, are capable of effectively capturing the multi-level features of images, thereby facilitating groundbreaking advances in tasks such as image classification, object detection, and image segmentation.

The application of CNN to the problem of RFI mitigation is, in essence, treating RFI as a specialized case within the broader domain of image processing. By capitalising on the robust capabilities of CNN in recognising and distinguishing patterns, RFI mitigation can be conceptualised as an image processing task. The time-frequency spectrum, where time and frequency serve as the two axes and signal amplitude corresponds to pixel values, allows radio frequency data to be transformed into an image-like format. This transformation enables CNN to be effectively trained to identify and mitigate RFI within this data.

This chapter will examine the ways in which the RFI mitigation problem can be formulated as an image processing task, and will investigate the potential of CNN as a means of addressing this challenge. Firstly, we will discuss approaches that frame the RFI mitigation problem as an image segmentation task in Section 6.1. Secondly, the potential to formulate RFI mitigation problem as an approach to the image restoration problem will be explored in Section 6.2. Finally, our proposed approach will be introduced in Section 6.3, a deep neural network based on image restoration.

6.1 Formulating RFI mitigation as image segmentation

The concept of formulating the RFI mitigation problem as an image segmentation problem is inspired by the labeling strategy that is commonly employed in traditional RFI removal methods. In the conventional approach, RFI is typically identified and localized by marking potentially disruptive regions using manual or automated techniques. These marked regions can then be targeted for subsequent interference elimination or mitigation. By treating RFI as a distinct region within an image, we can transform the problem into an image segmentation task, enabling more precise and automated processing, which enhances the overall accuracy and efficiency of the RFI mitigation process. Most of the work presented in the Section 3.4.2 uses this transformation strategy, with the most representative work using U-Net (Akeret et al., 2017a) and RFI-Net Yang et al. (2020).

In image segmentation, the objective is to assign a class label to each pixel within an image, rather than applying a single label to the image as a whole. This concept can be directly applied to the problem of RFI mitigation, where the time-frequency spectrum is considered as the input image. Each pixel in this spectrum corresponds to a specific time-frequency bin, representing the signal intensity at that particular moment and frequency. The goal of the RFI flagging problem is to classify each pixel in the spectrum as either "contaminated by RFI" or "non-contaminated" (i.e., normal astronomical signal). This can be mathematically formulated as a binary classification problem at the pixel level.

Transforming the RFI mitigation issue into an image segmentation task fundamentally involves categorizing each pixel within the time-frequency spectrum to determine whether it belongs to an RFI-contaminated region.This transformation allows us to leverage CNN to automatically detect and label these contaminated pixels. RFI typically manifests as distinctive patterns in the spectrum, such as

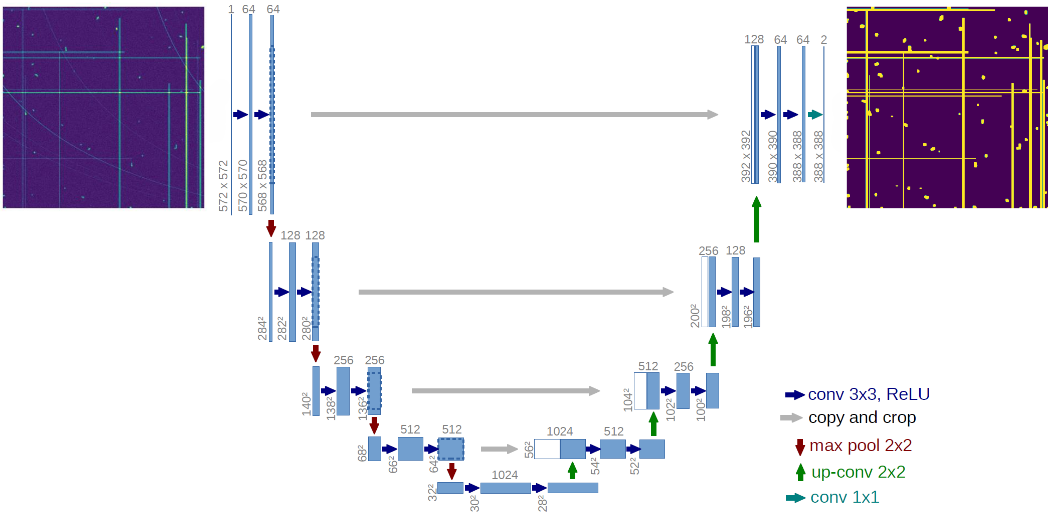

Figure 6.1: Illustration of U-Net applied for RFI mitigation. The input image is a RFI-corrupted dynamic spectrum and output is RFI mask provided by U-Net. The U-Net architecture in the figure is extracted from Ronneberger et al. (2015).

strong linear or block-like interference signals, making it well-suited for a pixel-level binary classification approach.

For RFI mitigation problem, the input time-frequency spectrum is denoted as $\mathbf{X}$, where $X_{i,j}$ represents the pixel value at position (i, j) in the spectrum. For each pixel $x_{i,j}$, we want to predict its category $y_{i,j}$:

$$y_{i,j} = \begin{cases} 1 & \text{if } x_{i,j} \text{ belongs to be in the RFI-contaminated region} \\ 0 & \text{if } x_{i,j} \text{ belongs to be in the non-contaminated region} \end{cases} \tag{6.1}$$

Following the formulation of model (5.30) on the image segmentaion problem, the mathematical expression for the RFI mitigation problem can be expressed in a similar manner.

$$\hat{\mathbf{Y}} = f(\mathbf{X}, \theta) \tag{6.2}$$

Here, $\mathbf{Y}$ is the predicted segmentation map, with each element $\hat{y}_{i,j}$ corresponding to the predicted class label for the pixel at position (i, j) in the input spectrum $\mathbf{X}$, and $f(\mathbf{X}, \theta)$ is the function representing the CNN. Typically, there is an output layer (also known as the last layer) that provides a probability value $p_{i,j}$ for each pixel (i, j), indicating the likelihood that the pixel belongs to a specific class (for example, RFI-contaminated). This probability value is calculated by the preceding layers of the model, and then a threshold τ is applied to make the final classification decision (for binary classification, threshold is set to 0.5). This process can be described by the following mathematical expression:

$$\hat{y}_{i,j} = \begin{cases} 1 & \text{if} \quad p_{i,j} \geq \tau \\ 0 & \text{if} \quad p_{i,j} < \tau \end{cases} \tag{6.3}$$

A commonly used architecture for this type of segmentation task is the U-Net, which consists of an encoder-decoder structure. The encoder compresses the input image into a lower-dimensional representation, while the decoder reconstructs the segmentation map from this representation, outputting pixel-level predictions. A practical example of this approach is a study that employs a U-Net architecture to perform RFI labeling (Akeret et al., 2017a) . In this work, the time-frequency spectrum were fed into the U-Net, which was trained to output a binary segmentation map for each spectrum. Figure 6.1 provides an illustration of the U-Net architecture, highlighting not only the network structure but also the flow of time-frequency data through the network. The U-Net successfully learned to differentiate between RFI-contaminated and non-contaminated regions, achieving accurate pixel-level classification.

6.2 Formulating RFI mitigation as image restoration

As discussed in Section 3.4, most existing methods in the literature approach the RFI mitigation problem as a detection or labelling task. Therefore, we can translate the RFI mitigation problem into image segmentation problem. These methods focus primarily on localizing the time-frequency bins in the dynamic spectra that are likely affected by RFI, distinguishing between RFI-corrupted and RFI-free bins. These methods mitigate the effects of RFI by removing the flagged portions of the data in subsequent processing steps. While this approach does reduce the impact of RFI to some extent, it also leads to the inevitable loss of valuable astronomical data, focusing more on RFI detection than actual removal. In contrast, we propose a more challenging approach by framing RFI mitigation as a restoration task. This shift in perspective introduces several significant challenges, the most prominent being the stringent requirements for the dataset. Unlike the simpler segmentation task, where the dataset only requires binary labels (0 for non-RFI and 1 for RFI-contaminated regions), the restoration task demands access to pristine, RFI-free astronomical data. Such data is crucial because the model must learn not only to detect but also to reconstruct the underlying astronomical signals that have been distorted by RFI.

To address this challenge, we leverage the simulated data generation framework introduced in the Chapter 4. This framework is specifically designed to produce high-quality datasets that include both clean astronomical data—such as pulsar signals—and artificially generated RFI signals. By combining these elements, the framework can simulate realistic, RFI-contaminated dynamic spectra while still retaining the original clean signals required for model training. This capability makes it feasible to reframe the RFI mitigation problem as a restoration task, allowing us to develop and evaluate advanced machine learning models for recovering the underlying astronomical signals from corrupted data.

In the meanwhile, our main rationale is that recent advances in machine learning, particularly in deep learning and CNN, offer powerful tools for tackling complex image restoration tasks. These advancements have significantly improved the ability to recover clean images from degraded or noisy ones. These developments are directly applicable to the RFI mitigation problem. By leveraging state-of-the-art CNN, we can directly recover clear, RFI-free dynamic spectra from measurements that have been corrupted by RFI. The ability of these networks to capture intricate patterns and dependencies within data makes them particularly well-suited for identifying and removing the structured noise associated with RFI, thereby preserving the integrity of the underlying astronomical signals.

Image denoising, as a specific type of image restoration, involves recovering a clean image $\mathbf{X}$ from a degraded image $\mathbf{Y}$ that has been corrupted by a particular type of noise. When this noise is considered to be additive, denoted by $\tilde{\mathbf{E}}$, the relationship between the clean and corrupted images can be expressed by the following model:

$$\mathbf{Y} = \mathbf{X} + \tilde{\mathbf{E}}. \tag{6.4}$$

The denoising model (6.4) can be linked to the observation model (4.1) in the simulation framework, and by comparing the two models, a relationship between the two can be established. In the observation model (4.1), measured dynamic spectra $\mathbf{S}$ (resp. RFI-free signal $\mathbf{P} + \mathbf{E}$) can be associated with the observed image $\mathbf{Y}$ (resp. clean image $\mathbf{X}$) in the formulation of denoising task. Moreover, the RFI signal $\mathbf{R}$ in the observation model (4.1) can be interpreted as the noise term in the context of denoising model (6.4). This connection between the two tasks enables us to approach RFI mitigation as an image denoising problem.

6.3 Proposed approach: RFI-DRUNet

Image denoising is a quintessential image-to-image translation task, effectively tackled by deep networks with encoder-decoder architectures, which consistently yield outstanding results. In such architectures, the encoder extracts features from the input image at various levels while compressing the data, and the decoder reconstructs the image using these features through skip connections. One of the benefits of translating RFI mitigation into imgae denoising is that we can easily take advantage of recently proposed deep convolutional neural networks specifically designed for image denoising tasks, such as Zhang et al. (2017) and Zhang et al. (2020).

In this work, we have tailored the well-known DRUNet network (Zhang et al., 2021) to suit the specific needs of the RFI mitigation task. DRUNet is a deep convolutional network based on an encoder-decoder structure, enhanced by a residual module that significantly improves feature extraction capabilities. Traditional DRUNet models can handle varying noise levels in data by jointly processing the noisy image and a corresponding noise map as inputs during training. However, for the RFI mitigation task at hand, where interference in dynamic spectra can vary in intensity, this level of flexibility is unnecessary. Consequently, the RFI-DRUNet is designed to take only RFI-contaminated dynamic spectra as inputs during training, omitting the need for a noise map.

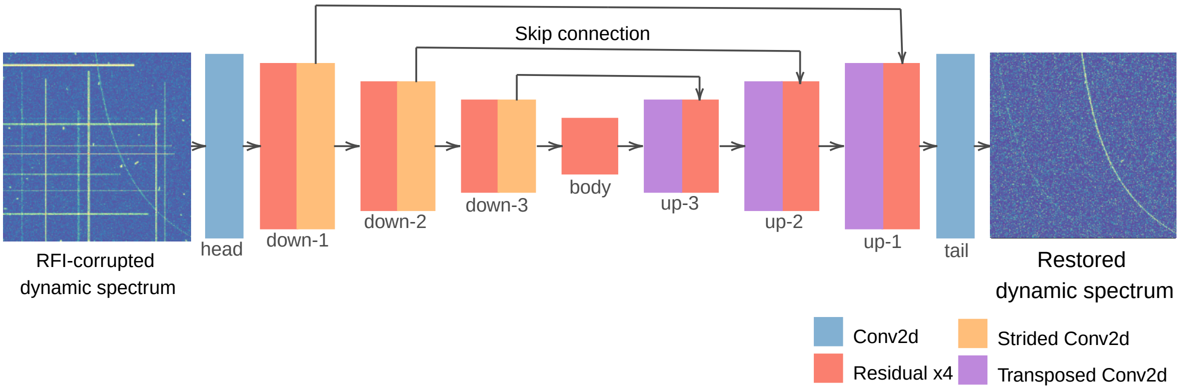

Figure 6.2: Architecture of the proposed RFI-DRUNet network. It takes as inputs RFI-corrupted dynamic spectra and provides as output restored (i.e., RFI-free) dynamic spectra. Details about the layers are provided in Table 6.1.

6.3.1 Architecture of RFI-DRUNet

The proposed RFI-DRUNet network, illustrated in Fig. 6.2, is structured around three key components: the encoder, the bottleneck, and the decoder, each contributing to the process of feature extraction, dimensionality reduction, and image reconstruction, respectively.

At the outset, the network begins with a head module that performs a 2D convolution operation with a 3×3 kernel, resulting in an output of 64 channels from an input size of $64 \times 64 \times 1$, where the first two dimensions represent the height and width, and the last dimension represents the number of channels. This convolution is conducted with a stride of 1×1 and padding of 1×1, ensuring that the spatial dimensions of the input are preserved.

Following the head module, the network progresses through three downsampling stages, each consisting of a combination of residual blocks and strided convolution. Each residual block is composed of a sequence of operations: a convolutional layer, followed by a ReLU activation, and then another convolutional layer. In these stages, the first residual block operates with a $3 \times 3 \times 64$ kernel, maintaining the input size of $64 \times 64 \times 64$. A strided convolution layer follows, with a 2×2 kernel and stride of 2×2 , doubling the channel depth to 128 and halving the spatial dimensions to 32×32. This pattern is repeated in the subsequent downsampling stages, with the second stage using a $3 \times 3 \times 128$ kernel and a 2×2 strided convolution that increases the channel depth to 256, reducing the spatial dimensions to 16×16. The third stage further deepens the network with a $3 \times 3 \times 256$ kernel and a strided convolution that produces an output size of $8 \times 8 \times 512$.

At the core of the network lies the bottleneck, which consists of four residual blocks utilizing a $3 \times 3 \times 512$ kernel. This layer maintains the spatial dimensions while further processing the data through deep convolutional operations. The upsampling process, which mirrors the downsampling structure, begins by reversing the dimensionality reduction with transposed convolutions. These layers use a 3×3 kernel with a stride of 2×2 to upsample the spatial dimensions while halving the channel depth. The first upsampling stage transforms the input from $8 \times 8 \times 512$ to $16 \times 16 \times 256$, with subsequent stages continuing this pattern, culminating in an output of $64 \times 64 \times 64$ at the final upsampling stage. Each upsampling stage also includes a set of four residual blocks, identical in structure to those in the downsampling stages. The network concludes with a tail layer, a simple convolution with a 3×3 kernel that reduces the channel depth to 1, producing an output size of $64 \times 64 \times 1$.

Overall, the RFI-DRUNet architecture efficiently captures and reconstructs the input data using a combination of deep residual learning and convolutional operations. This design enables the network

Module	Block	Operation	Kernel size	Stride	Padding	Input size	Output Size
head	Conv2d	Conv2d	3 × 3 × 64	(1, 1)	(1, 1)	(64, 64, 1)	(64, 64, 64)
down-1	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 64	(1, 1)	(1, 1)	(64, 64, 64)	(64, 64, 64)
	Strided Conv2d	Conv2d	2 × 2 × 128	(2, 2)	(1, 1)	(64, 64, 64)	(32, 32, 128)
down-2	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 128	(1, 1)	(1, 1)	(32, 32, 128)	(32, 32, 128)
	Strided Conv2d	Conv2d	2 × 2 × 256	(2, 2)	(1, 1)	(32, 32, 128)	(16, 16, 256)
down-3	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 256	(1, 1)	(1, 1)	(16, 16, 256)	(16, 16, 256)
	Strided Conv2d	Conv2d	2 × 2 × 512	(2, 2)	(1, 1)	(16, 16, 256)	(8, 8, 512)
bottleneck	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 512	(1, 1)	(1, 1)	(8, 8, 512)	(8, 8, 512)
up-3	Transposed Conv2d	Conv2d	2 × 2 × 256	(2, 2)	(0, 0)	(8, 8, 512)	(16, 16, 256)
	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 256	(1, 1)	(1, 1)	(16, 16, 256)	(16, 16, 256)
up-2	Transposed Conv2d	Conv2d	2 × 2 × 128	(2, 2)	(0, 0)	(16, 16, 256)	(32, 32, 128)
	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 128	(1, 1)	(1, 1)	(32, 32, 128)	(32, 32, 128)
up-1	Transposed Conv2d	Conv2d	2 × 2 × 64	(2, 2)	(0, 0)	(32, 32, 128)	(64, 64, 64)
	Residual ×4	(Conv2d+ReLu+Conv2d) ×4	3 × 3 × 64	(1, 1)	(1, 1)	(64, 64, 64)	(64, 64, 64)
tail	Conv2d	Conv2d	3 × 3 × 1	(1, 1)	(1, 1)	(64, 64, 1)	(64, 64, 1)

Table 6.1: Details of the layers of the proposed RFI-DRUNet network. The size of input data and output data are ($H \times W \times C$), where H, W and C stand for height, width and the number of channels, respectively. Sizes of the input can vary depending on the needs; an input data size of 64×64 is chosen here as an example.

to effectively mitigate RFI in dynamic spectra by learning complex patterns and removing unwanted interference, ensuring high quality outputs suitable for further analysis.

6.3.2 Core building block of the RFI-DRUNet architecture

After presenting the architectural details of the RFI-DRUNet, some of the core mechanisms of the network are described in this section. These mechanisms are key to the performance and efficiency of the RFI-DRUNet, and together they enable the network to excel in handling complex RFI-interfered data.

Residual connection

As the number of network layers increases, the problem of gradient vanishing or gradient explosion can seriously affect the training process. To overcome this problem, ResNet (He et al., 2016) was designed with residual connection that allows the network to learn residual mappings, i.e., to add inputs directly to the layer's outputs, thus facilitating the flow of gradients. The fundamental concept behind residual

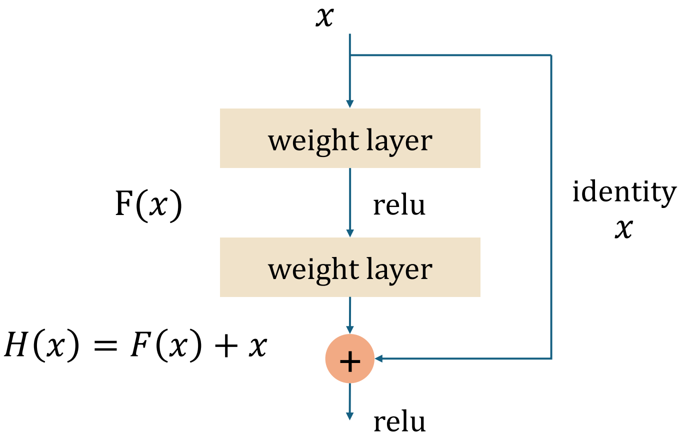

Figure 6.3: Illustration of residual connection a CNN.

connections is to frame the learning task as one of learning a residual mapping instead of directly learning an unreferenced mapping. This residual connection in a neural network is depicted in the Figure 6.3.

In mathematical terms, $H(x)$ is the desired underlying mapping to be learned by a few stacked layers, and $F(x)$ is the actual mapping provided by these layers, then instead of learning $H(x)$ directly, the network learns a residual function $F(x) = H(x) - x$. The output of the stacked layers is then given by $H(x) = F(x) + x$, where x is input to the layer.

In the proposed RFI-DRUNet network, extensive use of residual connections is applied, where each module of the network contains four stacked residual layers, except for the head and tail, which are designed to improve the learning capability of the network.

Downsampling operation

Downsampling is a process of reducing the spatial dimensions of data; in image-to-image applications, downsampling reduces the spatial dimensions of the feature maps, which not only reduces the amount of computation and memory consumption, but also increases the receptive field of the network, allowing the network to extract various features such as edge information locally and global information of the overall shape as the spatial resolution is progressively reduced. Popular downsampling methods include various pooling operations (maximum pooling and average pooling) and convolution operations with a stride longer than one.

Upsampling operation

Upsampling is a process of transforming data from low resolution to high resolution, in the context of image segmentation or image denoising, its necessary to restore the feature map to the resolution of the original image in order to facilitate pixel level operations. The main challenge in upsampling is to reconstruct the spatial information loss in the downsampling process while maintaining the reliability of the data. In addition, it is necessary to consider computational efficiency in terms of both spatial and temporal complexity. Common upsampling techniques include nearest-neighbor interpolation, bilinear interpolation, Transposed convolution, and pixel shuffle, among others. In the RFI-DRUNet, Transposed convolution has been adopted mainly because it can be trained to obtain features that recognize and reconstruct complex images better than traditional interpolation methods. The downsampling operation used in RFI-DRUNet is a convolutional layer with a step size of 2. This is due to the fact that the parameters of the convolutional layer can be learned, and this learning ability can be more flexible to adapt to the different noise patterns of the image when the task of image denoising is heavy.

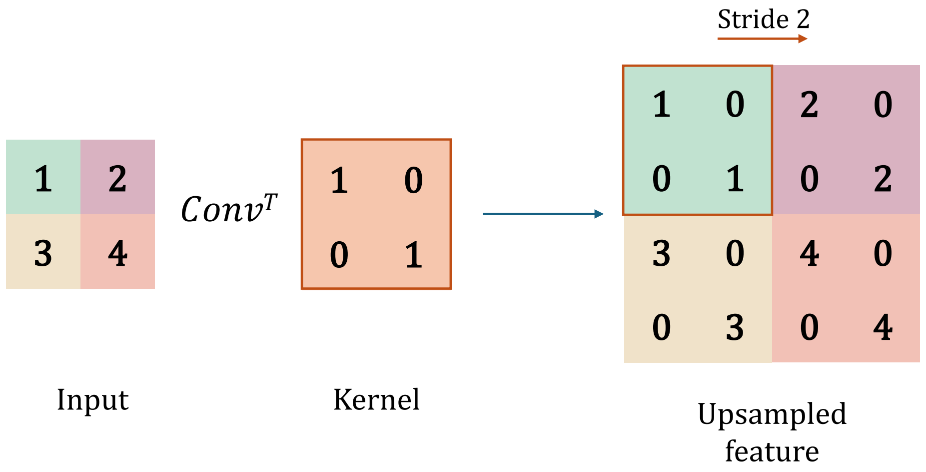

Figure 6.4: Demonstration of the transposed convolution operation in a CNN. A 2 × 2 input matrix is convolved with a 2 × 2 kernel using a stride of 2, resulting in a 4 × 4 upsampled output.

Transposed convolution, also referred to as deconvolution, is essentially the counterpart to the traditional convolutional operation outlined in Section 5.3.1. In traditional convolution, the input data is slid through a convolution kernel to produce a smaller-sized output feature map, whereas transposed convolution produces a larger-sized output. In transposed convolution, the first step is to spatially expand the input feature map, usually by circumferentially filling the map with zeros. This spatial expansion is essential for subsequent convolutional operations. Convolutional kernels are then used to perform sliding window operations over the expanded feature map. Unlike standard convolutions, the output of each kernel in a transposed convolution includes a larger region of the output feature map, effectively scaling up its dimensions. The final dimensions of the output feature map can be fine-tuned by adjusting the convolution step, with higher steps resulting in more pronounced expansions. This feature gives transposed convolution the flexibility to fine-tune the degree of upsampling. A key advantage of transposed convolution is its ability to learn, allowing it to skilfully reconstruct high-resolution detail from initially low-resolution feature maps through the training process. This is particularly important for applications such as image super-resolution and semantic segmentation. Figure 6.4 illustrates feature maps are upsampled through the application of transposed convolution.

The calculation between the input size and output size of the transposed convolution can be expressed as follows,which is the opposite of Eq (5.28):

$$\text{Output size} = (\text{Input size} - 1) \times \text{Stride} + \text{Kernel size} - 2 \times \text{Padding} \tag{6.5}$$

Résumé du chapitre

Dans ce chapitre, nous avons étudié l'utilisation de réseaux neuronaux convolutifs (CNN) pour la réduction des RFI. Les CNN ont démontré leur efficacité dans le domaine du traitement d'image, en capturant les caractéristiques d'image multi-échelles et en contribuant à des progrès significatifs dans des domaines tels que la classification d'image, la détection d'objet et la segmentation d'image. L'application des CNN à la réduction des RFI est essentiellement une application spécialisée du traitement d'image. En considérant le temps et la fréquence comme des axes, et l'amplitude du signal comme les valeurs des pixels, les données radiofréquences peuvent être mises en forme de manière similaire à une image. Cela permet aux CNN d'être entraînés efficacement pour détecter et réduire les RFI dans les données.

Nous avons d'abord abordé la réduction des RFI en le considérant comme une tâche de segmentation d'image. Les RFI sont traditionnellement identifiées et localisées en marquant les régions potentiellement affectées par des RFI à l'aide de techniques manuelles ou automatiques. Les données dans ces régions identifiées subissent ensuite un processus de suppression ou d'atténuation des interférences. En considérant les RFI comme une région distincte à l'intérieur d'une image, nous pouvons donc transformer le problème en une tâche de segmentation d'image, ce qui permet un traitement plus précis et automatisé. Cela améliore l'exactitude globale et l'efficacité du processus de réduction des RFI. Dans le domaine de la segmentation d'image, l'objectif est d'attribuer une étiquette de classe à chaque pixel de l'image, plutôt que d'appliquer une seule étiquette à l'image entière. Cette approche est directement applicable au problème de réduction des RFI, où le spectre dynamique est traité comme l'image en entrée. Chaque pixel de ce spectre correspond à case temps-fréquence spécifique, représentant l'intensité du signal à un moment et à une fréquence donnés. L'objectif de la détection des RFI est de catégoriser chaque pixel du spectre comme étant soit “contaminé par des RFI”, soit “non contaminé” (c'est-à-dire un signal astronomique normal). Cela peut être formulé mathématiquement comme un problème de classification binaire au niveau du pixel.

Enfin, nous avons présenté une approche de réseau neuronal profond basée sur la restauration d'image. Le débruitage d'image est une tâche classique de restauration, abordée efficacement par des réseaux profonds avec des architectures d'encodeur-décodeur, qui ont produit des résultats exceptionnels. Dans ces architectures, l'encodeur extrait des caractéristiques de l'image d'entrée à plusieurs niveaux tout en compressant les données, et le décodeur reconstitue l'image à partir de ces caractéristiques à l'aide de connexions de liaison. L'avantage de considérer la réduction des RFI comme un problème de débruitage d'image est que nous pouvons bénéficier des performances obtenues par des réseaux neuronaux convolutifs profonds récemment développés, spécifiquement conçus pour les tâches de débruitage. Dans cette thèse, nous avons adapté le réseau DRUNet, bien établi, pour répondre aux besoins spécifiques de la tâche de réduction des RFI.

Chapter 7

Experimental framework and results

7.1 Experimental framework

In order to evaluate the performance of the proposed RFI-DRUNet method, we plan to conduct a series of extensive and in-depth experiments. These experiments aim to demonstrate the effectiveness of the method in recovering spectrum corrupted by RFI. By comparing the recovery results under different conditions, we will show how RFI-DRUNet can successfully reconstruct high-quality spectra that closely resemble the original signal, even when dealing with various intensities and types of RFI interference. By comparing the recovery results under different conditions, we will show how RFI-DRUNet can successfully reconstruct high quality spectra that closely resemble the original signal, even when dealing with different intensities and types of RFI interference. These experiments will not only cover typical interference scenarios, but will also include tests under extreme conditions, providing a thorough validation of the robustness and applicability of the method. In addition, we will compare our proposed method with popular image segmentation techniques currently used for RFI detection, with the aim of proving that our network, when adapted to the task of RFI detection, performs as well as or better than these established methods. Through the analysis of experimental data, we aim to demonstrate the exceptional performance of the RFI-DRUNet method in the task of RFI mitigation and highlight its potential in real-world applications. Furthermore, we will show how this method effectively eliminates the effects of RFI in practical scenarios, further validating its applicability and effectiveness in real-world use cases.

Firstly, the proposed network will be trained and validated on a simulated dataset established in the Section 4.5.3. This simulated data set consists of eight different cases, designed based on the presence or absence of noise and the type of RFI involved. A separate model will be trained for each of these eight cases, resulting in eight different models. Following the training phase, cross-validation will be performed within specific scenarios (according to our design, these scenarios can be divided into two groups: those with background noise and those without). This cross-validation will use a test dataset to assess the ability of each model to recover data affected by a particular type of RFI.

Next, experiments are performed on data sets with different noise backgrounds and out-of-distribution data to evaluate the robustness of the models in dealing with unfamiliar RFI and different noise conditions. This phase of the experiment is crucial to test the generalisation ability of the models in real-world conditions where the RFI types and noise levels may differ from those in the training data.

In the second part of the experiment, the restoration results will be applied to the task of RFI detection. We will compare the proposed method against several popular segmentation methods, focusing on their ability to accurately detect RFI. This comparison will help us understand how well the proposed network can be adapted for RFI detection tasks, even when it was originally designed for RFI removal.

Finally, we will conduct an in-depth study of the architecture of model through ablation experiments. By systematically modifying the network structure, we aim to identify which components contribute most to the model's recovery performance. In addition, we will explore ways to reduce the network parameters to improve training and inference speed while minimising memory usage without sacrificing performance. This research is also aimed at evaluating the model potential for future use in real-world applications where efficiency and resource utilisation are critical considerations.

7.1.1 Implementation details

For each of the 8 simulated scenarios, a individual RFI-DRUNet model was trained, resulting in 8 different instances of the proposed network. For all models, the training parameters are kept consistent. During training, the generated dynamic spectra from the dataset $\mathcal{S}$ are randomly cropped into patches of size 64×64. Data augmentation techniques, including flipping and rotating, are then applied to enhance model robustness. The loss function is defined using the ℓ_1-norm, and training is conducted using the Adam optimizer with a mini-batch size of 64. The learning rate is initially set to 10^{-4} and is halved every 100,000 iterations until it reaches 5×10^{-7}. Training requires approximately 40 hours to complete 10,000 epochs, utilizing an Nvidia RTX 3080 GPU and implemented in Pytorch. The selection of batch size and number of epochs is influenced by the technical constraints of the available computing resources. The choice of hyperparameters, including the number of training epochs and batch size, is limited by the available computational resources, while other training settings such as the loss function, the optimizer, and the augmentation strategy, are adopted based on the DRUNet training approach described in Zhang et al. (2021).

7.1.2 Quantitative metrics

Based on the experiments we need to perform as described above, the experiments to evaluate the model can be divided into two main areas, i.e. restoration and detection. For the restoration aspect, which focuses on image restoration and denoising, peak signal-to-noise ratio (PSNR) is one of the most widely used and recognised metrics. This metric measures the ratio between the maximum possible power of a signal and the power of corrupting noise. Higher PSNR values indicate better image quality and closer resemblance to the original signal. The PSNR is de fined as:

$$\mathrm{PSNR} = 10 \log_{10} \frac{\max \mathbf{X}^2}{\mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}})} \tag{7.1}$$

where $\max \mathbf{X}^2$ represents the squared maximum value of the RFI-free signal, and $\mathrm{MSE}(\mathbf{X}, \hat{\mathbf{X}})$ denotes the mean squared error between the ground truth RFI-free signal $\mathbf{X} = \mathbf{S} - \mathbf{R}$ and its restored counterpart $\hat{X}$, as estimated by the algorithm. Specifically, in the noise-free case ($\mathbf{E} = 0$, Scenario 1), the observation model simplifies to $\mathbf{S} = \mathbf{P} + \mathbf{R}$, making this metric a measure of the quality of the restored pulsar signal, i.e., $\hat{\mathbf{X}} = \hat{\mathbf{P}}$.

For the detection aspect, which can be framed as a binary classification problem, the effectiveness of the mitigation method can be assessed using standard classification metrics. This involves first constructing a confusion matrix, which summarizes the counts of correct and incorrect classifications based on the presence of a specific target. In the context of RFI mitigation, an RFI occurrence is considered a positive instance, while an RFI-free signal is treated as a negative instance. To construct the confusion matrix, begin by gathering the predictions made by the model and comparing them to the actual ground truth labels. The process involves categorizing each prediction into one of four categories: *i)* true positive (TP): the number of instances where the model correctly predicted RFI; *ii)* false positive (FP): the number of instances where the model incorrectly predicted RFI when it was actually no-RFI; *iii)* true negative (TN): the number of instances where the model correctly predicted no-RFI; *iv)* false negative (FN): the number of instances where the model incorrectly predicted no-RFI when it was actually RFI. The confusion matrix can be represented in a tabular format as shown in the Table 7.1. From this confusion matrix, standard metrics such as precision, recall, and F1 score can be derived to evaluate performance.

	Predicted RFI	**Predicted No-RFI**
Actual RFI	TP	FN
Actual No-RFI	FP	TN

Table 7.1: Confusion matrix for RFI detection. TP represents true positives, FP represents false positives, TN represents true negatives, and FN represents false negatives.

The values in this confusion matrix allow for the calculation of various performance metrics. For example, the true positive rate (TPR), also known as recall, is calculated as:

$$\mathrm{TPR} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}} \tag{7.2}$$

This metric reflects the percentage of actual RFI instances that are correctly identified. The false positive rate (FPR) is calculated as:

$$\mathrm{FPR} = \frac{\mathrm{FP}}{\mathrm{FP} + \mathrm{TN}} \tag{7.3}$$

This metric indicates the proportion of RFI-free instances that are incorrectly identified as RFI. The TPR ranges from 0 to 1, with a higher value indicating better performance in detecting RFI. Similarly, the FPR ranges from 0 to 1, with a lower value indicating fewer errors in classifying RFI-free signals.

After constructing the confusion matrix, the elements of the matrix can be used to calculate metrics that provide deeper insights into the classification performance of the model. Precision, as defined in Equation (7.4) reflects the proportion of correctly identified positive instances among all instances predicted as positive. Its value ranges from 0 to 1, where a precision of 1 indicates that every instance predicted as positive is indeed positive, demonstrating perfect accuracy in positive predictions. Conversely, a precision closer to 0 suggests that many of the predicted positives are incorrect, indicating poor model performance in distinguishing positive instances. A higher precision value is desirable, as it signifies fewer false positives and greater confidence in the model's positive predictions.

$$\mathrm{prec} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}. \tag{7.4}$$

Recall, as indicated by Equation (7.5), represents the percentage of actual RFI instances that are correctly identified. Its value ranges from 0 to 1, where a recall of 1 indicates that all actual RFI instances have been correctly identified, reflecting perfect sensitivity. A recall closer to 0 signifies that many RFI instances have been missed, indicating poor performance in detecting true positives. A higher recall value is preferable, as it indicates that the model is effectively capturing most of the actual RFI instances, thus minimizing false negatives and ensuring a more comprehensive detection of positives.

$$\mathrm{rec} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}. \tag{7.5}$$

As defined in Equation (7.6), the F1 score, calculated as the harmonic mean of precision and recall, balances the two metrics. This score is particularly useful for datasets with imbalanced classes, where one class significantly outnumbers the other, a situation likely encountered in moderately corrupted dynamic spectra.

$$\mathrm{F1} = \frac{2 \times \mathrm{prec} \times \mathrm{rec}}{\mathrm{prec} + \mathrm{rec}}. \tag{7.6}$$

The model performance was also assessed using the area under the precision-recall curve (AUPRC), which evaluates the overall discriminatory ability of the compared models. This metric is particularly useful in scenarios with imbalanced datasets, where negative instances greatly outnumber positive ones, as it emphasizes the model's capacity to accurately identify positive examples while maintaining overall accuracy. Additionally, the receiver operating characteristic (ROC) analysis was conducted to provide a comprehensive evaluation of the model performance in binary classification tasks (Metz, 1978). ROC curves plot the TPR against the FPR, with the area under the ROC curve (AUROC), also known as the c-statistic (Hastie et al., 2009), serving as a critical performance metric. A higher AUROC score indicates superior discriminatory power, reflecting the model's effectiveness in distinguishing between positive and negative instances. An AUROC score close to 1.0 denotes a near-perfect balance between sensitivity and specificity, whereas a score near 0.5 suggests performance equivalent to random chance. Consequently, a higher AUROC score signifies better overall performance, particularly in the context of imbalanced datasets. Such an evaluation strategy has also been adopted by other RFI detection works (Mesarcik et al., 2022).

7.2 Experimental results

7.2.1 Restoration results

In this section, the results of the experiments designed to evaluate the effectiveness of the proposed RFI-DRUNet method in restoring dynamic spectra corrupted by RFI will be presented. The simulated dataset used for training comprises two main scenarios based on the presence or absence of noise, denoted as S_1 for the noise-free scenario and S_2 for the noisy scenario, with each scenario further divided into four different cases depending on the combination of RFI types involved(A to D). As stated in the Section 7.1, a total of eight models will be trained, each corresponding to one of the eight simulated datasets. These models will be labeled based on the specific dataset they are trained on. Specifically, the models will be denoted as $\mathsf{S}_\square\mathsf{C}_\triangle$, where the indices $\square \in \{1,2\}$ and $\triangle \in \{\mathrm{A},\mathrm{B},\mathrm{C},\mathrm{D},\}$ represent the respective scenario and case. The restoration performances of the RFI-DRUNet models trained on specific cases, from case A to case D, are evaluated not only on the testing data set generated according to the same case but also on data sets corresponding to the other cases. This comprehensive evaluation approach allows for a more detailed understanding of how various types of RFI affect the model's performance and helps to identify potential limitations. Additionally, testing in both noisy and noise-free conditions provides deeper insights into the model's resilience and generalizability across different scenarios. In order to evaluate the performance gains of the proposed model from a restoration perspective, we will calculate the PSNR between the RFI-corrupted signal and the RFI-free signal. We also employ an RFI detector, referred to as Oracle, which can precisely identify RFI and replace the compromised time-frequency bands with zeros or random values. The PSNR obtained from the RFI-corrupted pulsar signal recovered by the Oracle detector represents the theoretical upper limit of RFI recovery achievable by current conventional RFI detection methods. By comparing our model with these two additional approaches, we can more clearly demonstrate the superiority of our method in terms of restoration performance.

Noise-free data set

The restoration results, detailing the average PSNR and standard deviations calculated across the test set for the four models trained under scenario S_1 are reported in the Table 7.2. The results indicate that the model $\mathsf{S}_1\mathsf{C}_{\mathrm{A}}$ demonstrates a significantly different performance compared to the other models, as it achieves high accuracy only when tested on the dataset $\mathsf{S}_1\mathsf{C}_{\mathrm{A}}$. This discrepancy can be attributed to the limited diversity of RFI types (restricted to nbt RFI) present in the training set.

Data set Model		$\mathsf{S}_1\mathsf{C}_{\mathrm{A}}$	$\mathsf{S}_1\mathsf{C}_{\mathrm{B}}$	$\mathsf{S}_1\mathsf{C}_{\mathrm{C}}$	$\mathsf{S}_1\mathsf{C}_{\mathrm{D}}$
Data		43.17 ±0.93	35.02 ±3.14	35.87 ±2.99	31.92 ±1.33
Oracle		59.27 ±6.37	54.43 ±6.11	54.24 ±5.61	51.94 ±5.51
RFI-DRUNet	$\mathsf{S}_1\mathsf{C}_{\mathrm{A}}$	70.58 ±10.86	56.59 ±10.04	59.75 ±8.96	53.82 ± 8.10
	$\mathsf{S}_1\mathsf{C}_{\mathrm{B}}$	73.33 ±8.70	72.38 ±8.24	71.74 ±8.30	70.80 ± 7.80
	$\mathsf{S}_1\mathsf{C}_{\mathrm{C}}$	72.65 ±9.00	70.94 ±8.32	71.45 ±8.65	69.85 ± 7.93
	$\mathsf{S}_1\mathsf{C}_{\mathrm{D}}$	72.09 ±9.23	71.44 ±8.87	71.31 ±8.92	70.72 ± 8.61

Table 7.2: Scenario 1: restoration performance in terms of average PSNR and standard deviations computed over the test data sets. Here, standard deviation calculated as the square root of the variance of the PSNR scores across the test data sets.

Figure 7.1 gives the restoration performance of $\mathsf{S}_1\mathsf{C}_{\mathrm{A}}$ on the four test data sets. Consequently, the model struggles to identify and correct other types of RFI effectively. In contrast, the models $\mathsf{S}_1\mathsf{C}_{\mathrm{B}}$ to $\mathsf{S}_1\mathsf{C}_{\mathrm{D}}$ exhibit strong performance not only on test sets corresponding to their respective training data but also show comparable restoration capabilities when applied to other cases. The training data sets for models $\mathsf{S}_1\mathsf{C}_{\mathrm{B}}$ and $\mathsf{S}_1\mathsf{C}_{\mathrm{D}}$ each include only one type of RFI in addition to pulse-like (nbt) RFI- either nbct RFI or bbt RFI. However, these models are still capable of effectively handling test datasets containing RFI types that they did not encounter during training, including the other type of RFI they were not trained on, and even combinations of both types. The robustness of these models can be explained by the similarities in the shapes of nbct RFI and bbt RFI, despite their differences in the direction of spreading. The use of data augmentation techniques, such as rotations and flips, during the training phase likely contributes to this robustness, enabling the models to generalize effectively to both nbct and bbt RFI even when only one of these types is present in the training set. This augmentation helps the models learn a broader range of features, making them more adaptable to various RFI patterns in the test data. The RFI-corrupted spectrum in $\mathsf{S}_1\mathsf{C}_\mathrm{D}$ and the the restored spectra provided by models $\mathsf{S}_1\mathsf{C}_\mathrm{B}$ to $\mathsf{S}_1\mathsf{C}_\mathrm{D}$ are presented in Figure 7.2.

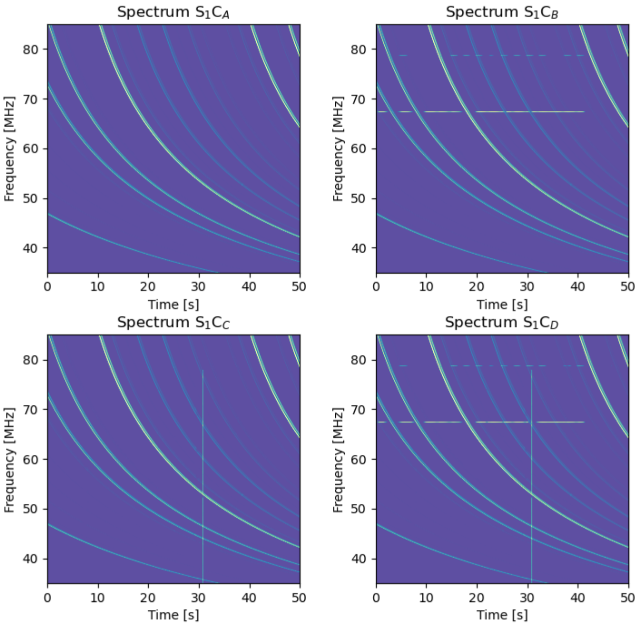

Figure 7.1: Visual comparison of spectra restored by $\mathsf{S}_1\mathsf{C}_\mathrm{A}$ across four test data sets in scenario S_1.

By comparing the PSNR values computed from the data itself (shown in the first row), it is evident that the four models, $\mathsf{S}_1\mathsf{C}_\mathrm{A}$ through $\mathsf{S}_1\mathsf{C}_\mathrm{D}$, achieve notable improvements in restoring RFI-corrupted spectra. Although the Oracle detector (replaced by zeros) also provides some level of restoration compared to the original data, its performance remains inferior to that of our proposed RFI-DRUNet model. Notably, even though model $\mathsf{S}_1\mathsf{C}_\mathrm{A}$ was trained solely on data sets containing pulse-like (nbt) RFI and was not exposed to nbct or bbt RFI during training, it still outperforms the Oracle detector on test datasets that include these untrained RFI types. This suggests that our proposed restoration model achieves superior performance in mitigating RFI effects compared to traditional methods, including the Oracle detector. The advantage of our approach is particularly evident in cases where the training data did not include the specific types of RFI encountered in the test sets, highlighting the robustness and effectiveness of our restoration technique.

Noisy data sets

Table 7.3 provides the results of the restoration in terms of average PSNR and standard deviations computed over the test set for the 4 models trained for the scenario S_2. Regarding the restoration performance of the four model produced by the proposed method, similar findings to those in S_1 can be found in the test results of S_2, except that the PSNR values of all test results in S_2 are lower than those in S_1 due to the presence of systematic noise. As with its performance in S_1, $\mathsf{S}_2\mathsf{C}_\mathrm{A}$ performs significantly better on the test data sets corresponding to its training dataset than on the other three datasets, the restoration results of $\mathsf{S}_2\mathsf{C}_\mathrm{A}$ over the four test data sets in S_2 are presented in Figure 7.3.

The performance of models $\mathsf{S}_2\mathsf{C}_\mathrm{B}$ to $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ on the four test data sets is consistent with the performance of their conterparts in S_1. In comparison to the data that has been corrupted by RFI, the restoration effect of the proposed method is markedly pronounced. It also exhibits a superior PSNR in comparison to

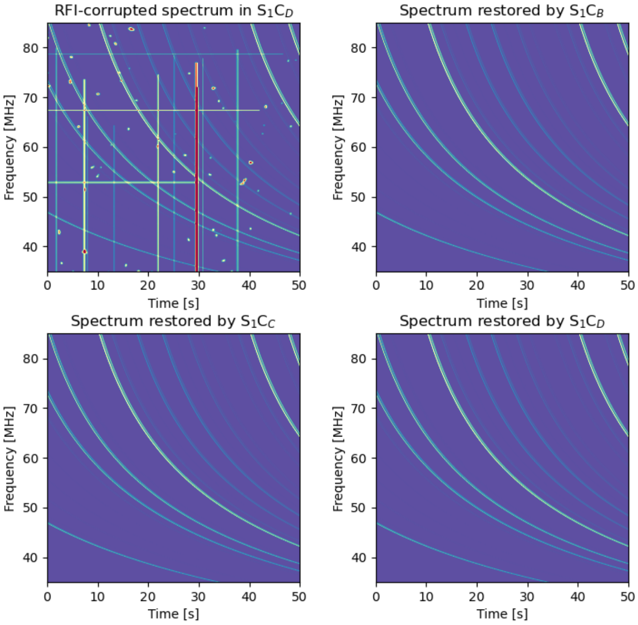

Figure 7.2: Visual comparison of RFI-corrupted spectrum in S_1C_D and its restored spectra provided by models S_1C_B to S_1C_D.

Model \ Data set		S_2C_A	S_2C_B	S_2C_C	S_2C_D
Data		43.17 ±0.93	35.02 ±3.14	35.87 ±2.99	31.92 ±1.33
Oracle		48.89 ±0.97	44.60 ±1.98	44.48 ±1.68	42.27 ±1.52
RFI-DRUNet	S_2C_A	59.43 ±6.54	36.38 ±4.88	37.77 ±4.27	32.89 ±1.69
	S_2C_B	60.15 ±5.08	59.81 ±4.91	59.88 ±5.02	59.22 ±4.73
	S_2C_C	60.36 ±4.34	60.01 ±4.40	60.12 ±4.27	59.51 ±4.25
	S_2C_D	60.41 ±4.49	60.18 ±4.55	60.19 ±4.47	59.95 ±4.53

Table 7.3: Scenario 2: restoration performance in terms of average PSNR and standard deviations computed over the test data sets.

the Oracle detector, which signifies an enhanced restoration effect. Figure 7.4 showcases the performance of the four model on the test data set S_2C_D

It is noteworthy that in the presence of background noise, the operation of the Oracle detector closely resembles the RFI mitigation techniques that are commonly used in real-world scenarios. Nevertheless, the RFI-DRUNet model, which we have proposed, consistently outperforms the oracle detector across all test datasets. This serves to illustrate the superior capability of our method in the restoration of data corrupted by RFI. The results demonstrate the robustness of our model in handling various types of RFI, even under noisy conditions, and highlight its clear advantage over traditional RFI mitigation approaches. This performance advantage demonstrates the efficacy of our approach in accurately recovering the underlying signal, which is pivotal for enhancing the reliability of data interpretation in practical applications. To visually illustrate the difference between our proposed RFI-DRUNet method and the

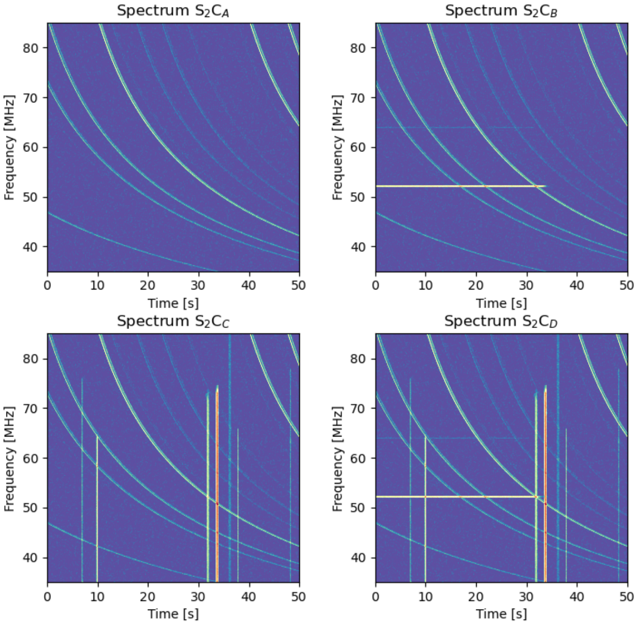

Figure 7.3: Visual comparison of spectra restored by $\mathsf{S_2C_A}$ across four test data sets in scenario $\mathsf{S_2}$.

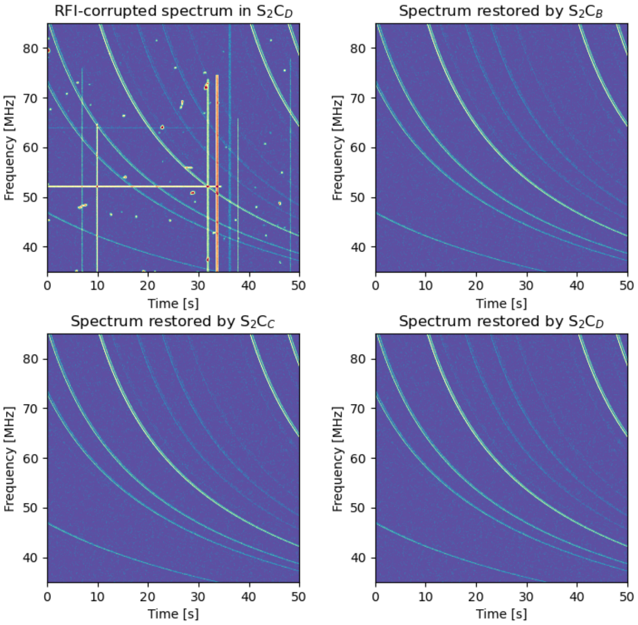

Figure 7.4: Visual comparison of RFI-corrupted spectrum in $\mathsf{S_2C_D}$ and its restored spectra provided by models $\mathsf{S_2C_B}$ to $\mathsf{S_2C_D}$.

Oracle detector, a comparative example will presented in the Figure 7.5 using spectra derived from test data sets $\mathsf{S_2C_D}$. These images will show the original RFI-corrupted signal, the signal processed by the Oracle detector, and the signal restored by our RFI-DRUNet model $\mathsf{S_2C_D}$. By comparing these visual representations, the superiority of our method in effectively removing RFI and accurately reconstructing the original signal will become evident. This example will clearly demonstrate how our approach outperforms the Oracle detector, particularly in scenarios where traditional RFI mitigation techniques fall short.

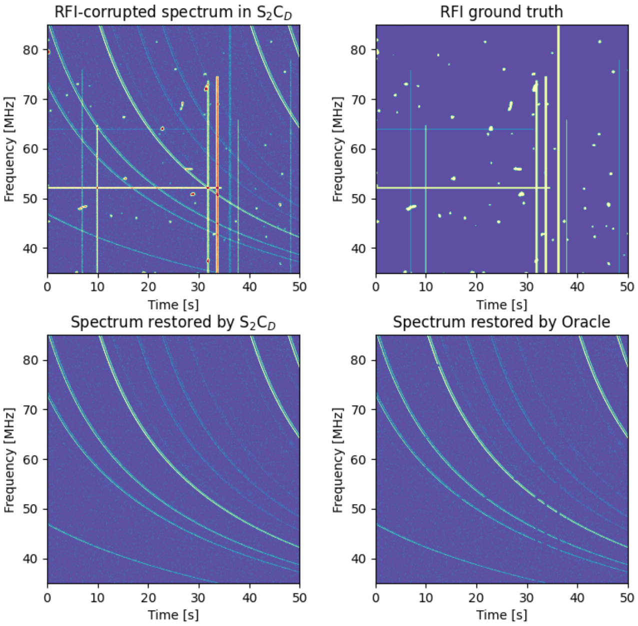

Figure 7.5: Visual comparison of RFI-corrupted spectrum in $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ and its restored spectra provided by model $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ and Oracle detector.

Validation

During the training process, it is essential to periodically evaluate the model on a separate dataset, known as the validation set, which is not used in the actual training. By monitoring validation metrics, such as PSNR (Peak Signal-to-Noise Ratio), we can assess whether the model is effectively learning. Typically, we expect the validation PSNR to increase as training progresses. If this trend is observed, it indicates that the model is successfully learning from the training data and generalizing well to the validation set. Moreover, the validation process plays a crucial role in helping us avoid overfitting, ensuring that the model does not merely memorize the training data but performs well on new, unseen data.

The Figure 7.6 illustrate the restoration performance, measured in terms of PSNR, as a function of the number of epochs during the validation for S_2. In S_2, the validation results improve with the increasing number of training epochs. The rate of improvement levels off around 4000 epochs, after which the restoration performance of each model on the validation set tends to converge. This figure demonstrates that the validation outcomes are consistent with those observed during the testing phase (S_2). In this scenario, the validation performance of the models is consistent with their restoration performance on the test datasets.

Given the performance of the various models, we will focus on the $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ model in the subsequent sections, unless otherwise specified. This model, trained with data that includes both background noise and all simulated types of RFI, is considered the most representative of real observational conditions.

Out-of-distribution data sets

In order to assess the resilience of the proposed RFI-DRUNet model in the context of experimental conditions that diverge from those encountered during the training phase, we now turn our attention to

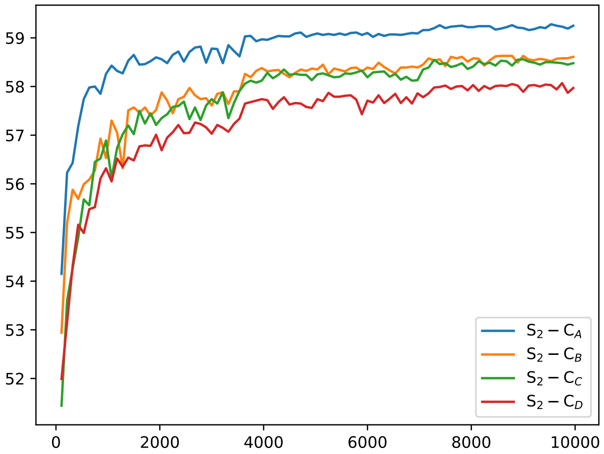

Figure 7.6: Validation: restoration performance (in terms of PSNR) as a function of the number of epochs.

examining the model's capacity for generalisation when analysing data that falls outside the distribution observed during training (out-of-distribution, or OOD data in hereafter). The term 'generalisation' is used to describe a model's capacity to maintain strong performance on unseen data. This is crucial for ensuring the reliability of the model in real-world applications. By investigating the model's ability to generalise with respect to OOD data, it is possible to gain insights into how the model performs when confronted with new environments, differing data characteristics and potential noise or interference. This hypothesis of model-data mismatch can be considered to have two primary aspects. Firstly, there is a discrepancy in background noise levels, which may vary significantly from those encountered during the training phase. Secondly, the occurrence of RFI types that were not observed during the training phase. The objective of this investigation is to gain a more comprehensive understanding of the RFI-DRUNet model's capability to generalise, ensuring that it can effectively handle a variety of unseen RFI conditions while maintaining high performance and stability in practical applications.

In regard to experiments involving supplementary background noise, it is important to note that during the generation of simulated data, the background noise was set with a variance of 1. This indicates that our proposed method was designed to restore signals corrupted by RFI under a background noise of $\sigma_{\mathrm{E}}^2 = 1$ (SNR ≈ -5dB). As we increase the noise level in the data, the background noise variance is adjusted to higher values such as $\sigma_{\mathrm{E}}^2 = 2$, 3, 5, and 10, with corresponding SNR $\in \{-8.6, -10.4, -12.6, -15.6\}$ [dB]. This adjustment represents increasingly severe noise conditions. Our evaluation aims to assess the performance of the model $\mathsf{S}_2\mathsf{C}_{\mathrm{D}}$ trained at the original noise level in the face of these more challenging noise environments. The Table 7.4 shows that as the noise level increases, the PSNR values decrease progressively. This indicates a deterioration in the model's restoration performance under higher noise levels. The decline can be attributed to the increased background noise, which masks more of the useful information in the signal, thereby making the signal restoration process more challenging. Despite this, the PSNR values remained above 50 dB even at higher noise levels, demonstrating that the model maintains effective restoration performance under challenging conditions. In practical observations, the background noise level can vary due to differing observational conditions. Therefore, the model must be capable of performing restoration tasks effectively across diverse noise environments. The results indicate that the model reliably operates under varying noise conditions, providing strong evidence of its robustness and effectiveness in real-world applications.

Since RFI in real observations may manifest in forms beyond the scope of our simulation framework, it is essential to evaluate the proposed method's ability to handle unseen types of RFI. Testing for the robustness of the method against out-of-distribution RFI is therefore a critical step in ensuring its practical applicability. In addition to the RFI profiles generated within our existing simulation framework, instances of RFI with sinusoidal shapes have been considered. Specifically, rather than generating RFI

	Simulation parameters		**PSNR**
OOD with respect to noise level	$\sigma_{\mathrm{E}}^2 = 1$	SNR $= -5.6$dB	59.95 ±4.53
	$\sigma_{\mathrm{E}}^2 = 2$	SNR $= -8.6$dB	59.24 ±0.81
	$\sigma_{\mathrm{E}}^2 = 3$	SNR $= -10.4$dB	57.81 ±0.61
	$\sigma_{\mathrm{E}}^2 = 5$	SNR $= -12.6$dB	54.86 ±1.56
	$\sigma_{\mathrm{E}}^2 = 10$	SNR $= -15.6$dB	50.84 ±2.15

Table 7.4: Out-of-distribution data sets: restoration performance in terms of average PSNR and standard deviations for different noise levels.

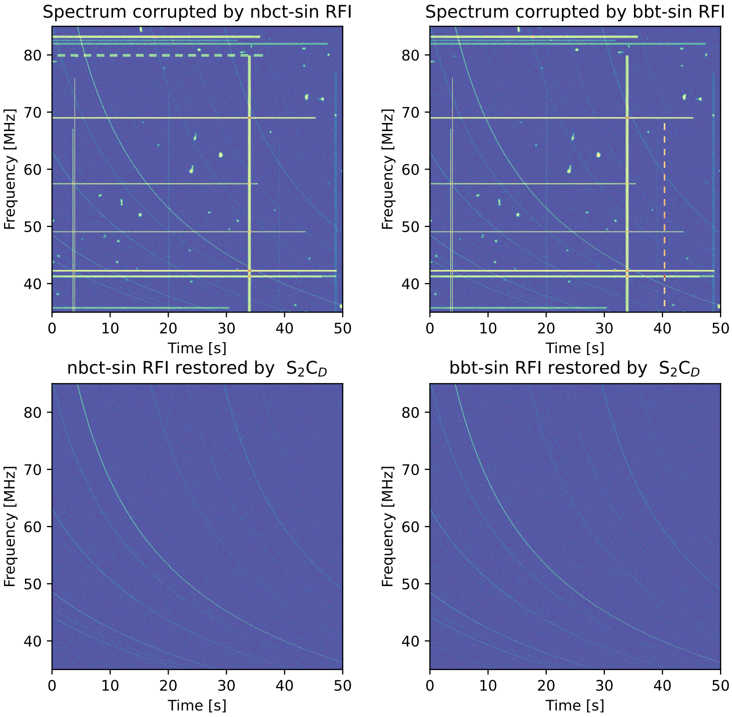

Figure 7.7: Visual comparison of spectra with sinusoidal RFI and their restored spectra provided by $\mathsf{S}_2\mathsf{C}_\mathrm{D}$, nbct-sin RFI with narrowband continuous temporal spreading (left top), and its restored spectrum (left bottom), bbt-sin RFI with broadband temporal and spectral spreading (right top), and its restored spectrum (right bottom).

signals solely with the Gaussian shapes described in Eq. (4.6), I introduce nbct and bbt RFI signals that follow an oscillating profile. These sinusoidal RFI profiles, denoted as nbct-sin and bbt-sin depending on their temporal and spectral spreading, are incorporated into the testing sets to assess the method's performance against more varied and complex RFI patterns. These two types of RFI, nbct-sin and bbt-sin, are illustrated in Figure 7.7.

The test data set associated with the model serves as the baseline, and the corresponding test sets are augmented with nbct-sin RFI, bbt-sin RFI, and a combination of both. This setup is used to evaluate the model's ability to handle previously unseen RFI types. Table 7.5 reported the $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ model's recovery performance on test sets that include these sinusoidal RFI instances. When nbct-sin RFI or bbt-sin RFI were added to the test datasets, the model's performance experienced a slight decrease compared to the baseline. However, this decline remains within a reasonable range, indicating that the model is still capable of effectively handling sinusoidal RFI.

In summary, it is clear that within a reasonable range of mismatches, the noise level has a negligible effect on the restoration performance. Similarly, the effect of sinusoidal RF interference in the test set on the restoration results is minimal, which further demonstrates the robustness of the proposed method.

	Simulation parameters	PSNR
OOD with respect to RFI profile	baseline	59.95 ±4.53
	baseline + nbct-sin RFI	59.90 ±3.22
	baseline + bbt-sin RFI	59.91 ±3.17
	baseline + nbct-sin RFI + bbt-sin RFI	59.59 ±3.11

Table 7.5: Out-of-distribution data sets: restoration performance in terms of average PSNR and standard deviations for different RFI profiles.

7.2.2 Detection results

Having demonstrated the effectiveness of our proposed RFI-DRUNet model in the restoration of RFI-corrupted signals, we now turn our attention to its performance in the task of RFI detection. In comparing the restoration results of our proposed method with those produced by the Oracle detector—an idealized detector with perfect knowledge of all RFI locations—we have demonstrated that our approach excels in recovering signals corrupted by RFI, surpassing the performance of current popular RFI detection methods. However, it is important to recognize that this comparison may not fully reflect the capabilities of these RFI detectors, as their underlying models were not specifically designed for the task of restoration. To provide a more comprehensive evaluation, we have adapted our proposed restoration method for use in RFI detection. This adaptation allows us to directly compare its performance with established RFI detection techniques, enabling a thorough assessment of the relative strengths and weaknesses of our approach in various signal processing tasks.

In order to achieve a comparison on the detection task, we transform the proposed method into a simple RFI detector. To do this, we need to construct a binary mask to indicate the presence of RFI by restoration results.

$$\hat{M}(n,k) = \begin{cases} 1, & \text{if } |\hat{S}(n,k) - S(n,k)| > \eta \quad \text{(presence of RFI)} \\ 0, & \text{otherwise} \qquad\qquad\qquad\quad\; \text{(absence of RFI)} \end{cases} \tag{7.7}$$

where $\hat{S}(\cdot,\cdot)$ represents the restored dynamic spectrum produced by RFI-DRUNet, $S(\cdot,\cdot)$ is the input dynamic spectrum, and η serves as a threshold that balances the trade-off between detection probability and false alarm probability. As done in previous studies by Kerrigan et al. (2019) and Mesarcik et al. (2022), the threshold is set at $\eta = 0.15$, which maximize the F1-score of RFI-DRUNet. The performance of the proposed method will be compared with the detection capability of two deep learning based RFI interference mitigation methods. The two networks selected for comparison, U-Net (Akeret et al., 2017a) and RFI-Net (Yang et al., 2020), have demonstrated outstanding performance in RFI detection tasks. Both networks are built upon classic deep learning architectures: U-Net, originally designed for medical image segmentation, has been successfully adapted for RFI mitigation , while RFI-Net leverages an encoder-decoder structure with residual blocks and batch normalization, making it a robust choice for RFI detection. These two models have been trained on the data sets corresponding to scenario $\mathsf{S}_2\mathsf{C}_\mathrm{D}$. Table 7.6 presents the detection performance comparison among the three models, evaluated using standard classification metrics such as precision, recall, F1-score, AUROC, and AUPRC. These metrics, introduced in Section 7.1.2, offer a comprehensive assessment of the binary classification capabilities of each model. The proposed RFI-DRUNet method exhibits exceptional performance, achieving a precision of 0.972 and a recall of 0.961, which translates to a high F1-score of 0.966. Additionally, RFI-DRUNet records the highest AUROC and AUPRC values among the three models, at 0.995 and 0.986, respectively. These metrics underscore RFI-DRUNet's excellence in both detecting and classifying RFI with notable accuracy and robustness. The relatively low standard deviations further indicate consistent performance across test datasets. Overall, these results establish RFI-DRUNet as the leading model for RFI detection, demonstrating superior precision, recall, and overall classification capability, particularly evident in its superior AUROC and AUPRC scores.

Figure 7.8 provides a qualitative comparison of the results from the different methods. The figure includes visualizations of the following: the original signal corrupted by RFI, the signal restored by RFI-DRUNet, RFI masks identified by U-Net and RFI-Net, the true RFI masks, and the RFI components removed by RFI-DRUNet. This comprehensive depiction allows for an in-depth assessment of each method's performance in handling RFI.

Method	**prec**	**rec**	**F1-score**	**AUROC**	**AUPRC**
RFI-DRUNet	0.972 ±0.398	0.961 ±0.0157	0.966 ±0.024	0.995 ±0.002	0.986 ±0.016
U-Net	0.858 ±0.048	0.987 ±0.005	0.917 ±0.029	0.991 ±0.003	0.926 ±0.027
RFI-Net	0.864 ±0.047	0.988 ±0.005	0.921 ±0.029	0.987 ±0.004	0.926 ±0.026

Table 7.6: Detection performance of compared algorithms in terms of precision, recall, F1-score, AUROC and AUPRC. The results are reported with mean and standard deviation computed over the test data sets.

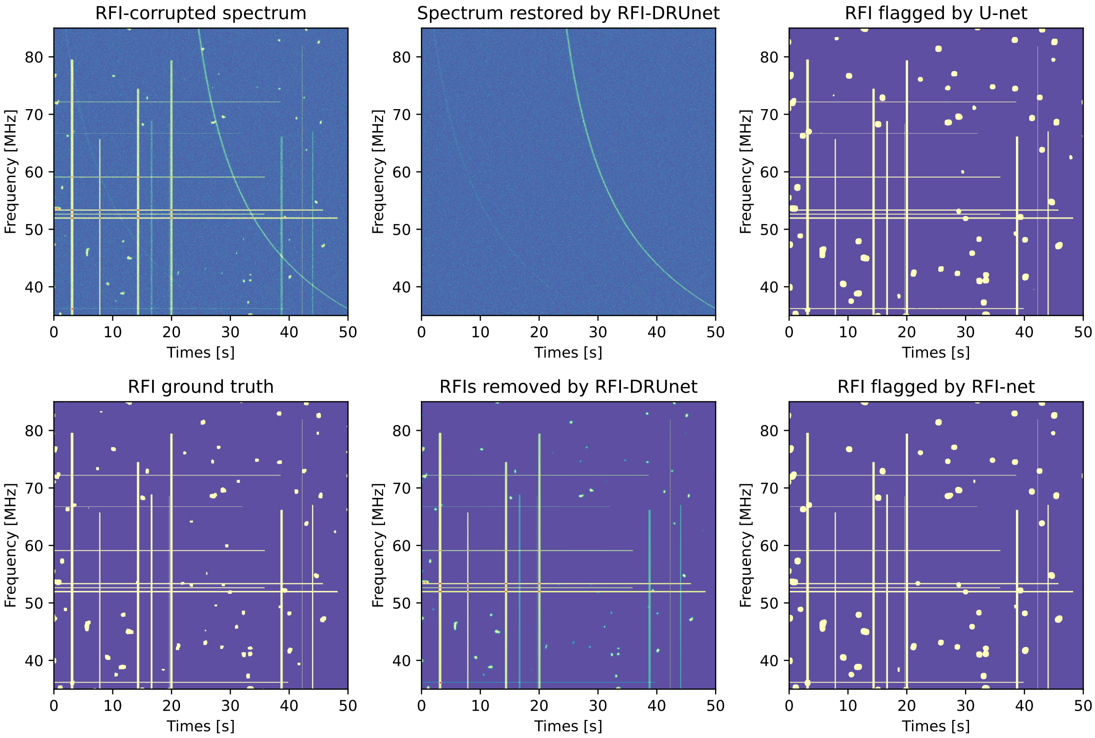

Figure 7.8: Visual comparison of the results for simulated spectrum from $\mathsf{S}_2\mathsf{C}_\mathsf{D}$ provided by the compared methods.

7.2.3 Ablation study

In order to achieve a more profound comprehension and to optimise the proposed RFI-DRUNet, as detailed in the Section 6.3, we conducted an ablation study. The objective of this ablation study was to quantify the contributions of key network components to the restoration performance, thereby identifying the elements that are crucial to the model's efficacy. Our focus was on two aspects that directly influence the model's complexity: the quantity of residual layers stacked within each module and the channel dimension throughout the network. By adjusting these parameters, we aim to find the optimal balance between complexity and accuracy, with the ultimate goal of simplifying the model without sacrificing performance. This approach is particularly relevant for deployment across diverse hardware environments.

In contrast to other ablation studies, which typically aim to assess the contribution of each module to model performance with the goal of optimising model performance and efficiency, our study specifically focuses on identifying ways to reduce model complexity and parameter count without compromising model effectiveness. By investigating the effects of varying the number of residual layers and channel dimensions, we aim to uncover the minimum configuration required to maintain high performance, thereby improving the model's suitability for use in resource-constrained environments.

The first concern in our ablation study is the number of residual layers in each downsampling and up sampling module. In the original RFI-DRUNet, four blocks of residual layers stacked in each module were used to increase networking capacity. For this purpose, we varied the number of residual blocks and denoted them by RL♢, where ♢ denotes the number of blocks. The performance of the model evaluated under configurations with RL4, RL3, RL2, RL1, correspond to 4, 3, 2, and 1 residual layers, respectively. We conducted a channel reduction experiment, testing the model with both half and quarter of the original channel number in the sampling modules to assess the impact on performance. In the original network, the channel configuration strata with 64 channels at the beginning of the downsampling module. As the network progresses through three downsampling modules, the number of channels doubles with each operation until it reaches 512. During three up sampling modules, the number of channels is halved with each operation, eventually returning to 64. This original channel configuration is referred to as ChOrig for ease of reference in subsequent discussions. The channel configuration is adjusted so that the number of channels is reduced to half and one-quarter of the original values. These configurations are denoted as ChR50 and ChR25 respectively, where R stands for "reduction," and the numbers indicate the percentage of the original channel count retained.

The experiments will examine the model's along two key dimensions: complexity and training speed. Model complexity promarily refers to the number of parameters and the floating-point operations per second (FLOPs) involved in the model. FLOPs measure the total number of floating-point arithmetic operations the model must perform to process a single input. Higher FLOPs typically indicate a more computationally demanding model, which can affect both the training and inference speed. The training speed is evaluated in terms of the time consumed by the network to complete one epoch training. The results of the experiments are reported in the Table 7.7 and the relationship between model complexity and its restoration performance are illustred in Figure 7.9.

RL	**CR**	**♯param.** ($\times 10^6$)	**FLOPs (G)**	**Time/Epoch (s)**	**PSNR (dB)**
RL4	ChOrig	32.65	2294.72	25.87	61.01
	ChR50	8.17	573.98	10.93	60.93
	ChR25	2.04	143.65	8.42	60.21
RL3	ChOrig	24.83	1753.55	20.48	60.98
	ChR50	6.21	438.69	9.46	59.57
	ChR25	1.55	109.82	7.89	59.37
RL2	ChOrig	17.01	1212.39	15.01	60.03
	ChR50	4.26	303.4	8.83	59.50
	ChR25	1.06	76.0	7.87	58.49
RL1	ChOrig	9.2	671.22	9.91	59.16
	ChR50	2.3	168.11	7.72	57.95
	ChR25	0.58	42.18	7.45	56.58

Table 7.7: Ablation study results showing the performance of different model's configurations in terms of parameter numbers (♯param.), FLOPs, time for training each epoch (Time/Epoch) and PSNR.

Based on the results of the ablation study, we observed that reducing the number of residual layers (RL) and the number of channels (CR) in the network effectively reduces the number of model parameters and the number of floating point operations (FLOPs), thereby reducing the complexity of the model. While this reduction in complexity reduces the time required to train each epoch, it also raises concerns about the potential weakening of the model's learning ability. Specifically, the model may fail to capture all of the complex features in the data, resulting in a decrease in PSNR values in image restoration tasks. In our experiments, we observed that a moderate reduction in the number of residual layers (RL), especially in deeper networks, can effectively decrease the complexity of the model without significantly affecting performance. Specifically, each reduction of one RL layer reduces the number of model parameters by approximately 25%. When moving from four RL layers to three, although the reduction in parameters affects the depth of the model, the impact on overall performance is relatively minimal. This may be because the capacity of the original model already exceeded the task requirements. However, when the number of RL layers is further reduced, especially to two or one, the performance drop becomes

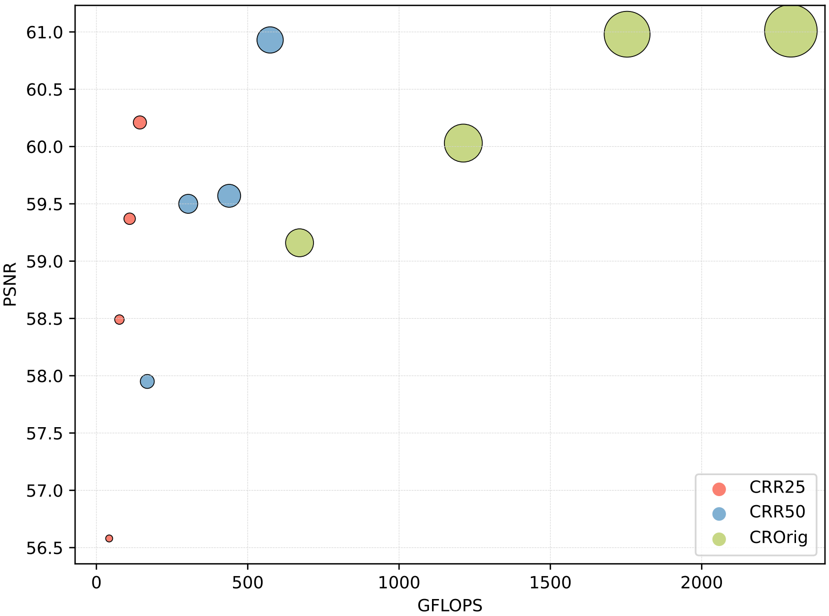

Figure 7.9: Illustrating the trade-off between model's complexity (FLOPs) and restoration performance in terms of PSNR. Each data point, represented by a circle with size proportional to the parameter count, and is colored to indicate its channel reduction configuration.

more pronounced, especially when the number of channels is also reduced. This more significant performance degradation can be attributed to the increased loss of hierarchical feature learning and feature description capabilities, preventing the model from effectively extracting features from the data.

In terms of reducing the number of channels, if we halve the number of channels in the network, the model parameters are correspondingly reduced by 75%. This directly reduces the complexity of the model and the computational requirements. Although the operation of reducing the number of channels can lead to a decrease in model performance, our experimental results indicate that this loss in performance is relatively small, especially compared to the effect of reducing the number of residual layers (RL). One explanation for this phenomenon is that the original network was designed for RGB images, which have three colour channels. However, our data contains only a single channel. Therefore, the original setting of the number of channels was inherently redundant for single-channel data and exceeded the needs for processing the single-channel data we provided. By reducing the number of channels, we were able to eliminate this redundancy without significantly affecting the model's ability to learn from and generalise the data.

These findings can help us better understand the relationship between model structure and performance, and provide valuable insights into how to adjust the model to maintain performance while reducing complexity.

7.2.4 Illustration on a real observation

In the previous Sections 7.2.1 and 7.2.2, proposed method has already demonstrated exceptional restoration and detection performance on simulated datasets generated according to the simulation framework presented in the Chapter 4. We now proceed to evaluate the method's efficacy using real observational data from the NenuFAR telescope. Owing to the absence of precise RFI annotations in these real datasets, quantitative metrics for model evaluation are challenging to apply. Consequently, we rely primarily on visual methods, comparing the restored dynamic spectra by our model with the original data to assess the effectiveness of RFI mitigation. To illustrate the performance of the model, we have selected a real observational dynamic spectrum from a pulsar observation of B1919+21. The selected spectrum spans a frequency range of 74 – 80 MHz and a duration of 5 seconds. Within this spectrum there is a notable presence of a nbct RFI around 78 MHz, and several clearly discernible pulsar signals are visible within the selected time frame, overlapping with the RFI. It is important to note that real data often exhibit complexities not found in the simulated data, possibly including interference types or noise patterns not considered by the framework. In the simulation framework, several simplifying assumptions have been made to facilitate data modeling, such as the assumption regarding the frequency variation of the pulsar signal. These simplifications may not fully align with real-world scenarios. Therefore, testing the model on real signals provides a robust assessment of its generalization capabilities to unknown data, which is crucial for verifying the model's practical applicability.

The models selected for comparison - RFI-DRUNet, U-Net and RFI-Net - were consistent with their settings in the Section 7.2.2. The comparative visualisation results are shown in Figure 7.10, where the top right panel shows the dynamic spectrum restored by RFI-DRUNet, while the bottom left and bottom right panels display the RFI detected by RFI-DRUNet and U-Net, respectively. RFI-Net's detection quality was deemed insufficient for comparison and is not shown here; its poor performance on real observations can be attributed to its limited generalisation and robustness capabilities. In the real dynamic spectrum, RFI manifests itself predominantly as a nbct RFI around 78 MHz, present briefly from 0.2s to 0.5s, then continuously from 1.5s to 3s, and then with a periodicity similar to our testing described in the Section ??. In addition, some RFI is present in the form of pulse-like (nbt).

Examining the restoration performance of RFI-DRUNet, it is evident that the model effectively recovers the spectrum corrupted by RFI, in particular the nbct RFI around 78 MHz. A closer examination of RFI-DRUNet's restoration around 78 MHz shows that the recovery around 77.5 MHz is quite successful, with the restored pulse signals clearly visible in this region. However, between 78 and 78.5 MHz, the restored spectrum shows a distinct distribution compared to other areas. RFI-DRUNet identifies the presence of RFI in this region, but the restoration is not entirely satisfactory, leaving traces of the RFI removal. The overall performance of RFI-DRUNet in restoring the spectrum is satisfactory; it visually removes RFI and restores the spectrum, although some areas show limitations, possibly due to the lack of training data featuring closely spaced, narrow nbct RFI in the 78-78.5MHz range, highlighting the generalisation limitations of the model.

The bottom two panels of the figure compare the RFI detection results between RFI-DRUNet and U-Net. It is clear that both RFI-DRUNet and U-Net can detect the majority of RFI, but RFI-DRUNet shows superior accuracy, detecting periodic RFI after 3 seconds and several instances of pulse-like RFI. In contrast, U-Net only detects nbct RFI between 0.2s and 3s, and even within this range it fails to identify all of the RFI present. These results confirm that RFI-DRUNet and U-Net can accurately detect RFI regardless of its shape and have better generalisation capabilities. It is also important to note that RFI-DRUNet not only restores the interfered dynamic spectrum, but also preserves most of the signals of interest.

7.2.5 Discussion

As previously mentioned in the Section 6.1, U-Net and RFI-Net effectively identified RFI locations in dynamic spectra by transforming the RFI mitigation problem into an image segmentation task. These methods use neural networks to extract features from two-dimensional data (typically images, in this case dynamic spectra) and perform pixel-level segmentation to detect RFI. While these approaches outperform RFI detection tasks, their limitation lies in the fact that they can only identify RFI signals, and the common strategy is to eliminate regions disturbed by RFI, which not only fails to restore the corrupted signals, but also results in the loss of the original astronomical data. Therefore, we proposed a novel approach that treats RFI mitigation as an image restoration task and proposed RFI-DRUNet, which not only accurately marks RFI but also reliably restores the damaged time-frequency segments. Another common strategy is to replace detected RFI locations with a preset reliable value, such as zero or a random number. In Section 7.2.1, we already discussed the strategy of replacing time-frequency segments corrupted by RFI with reliable values, using an oracle detector that knows all RFI locations to replace them with zero, and compared its recovery effect with that of our proposed RFI-DRUNet. Here we further analyse, utilizing both the U-Net and RFI-Net detectors, how these regions are filled with random values based on the instrument noise statistical model ($\mathcal{N}(0, \sigma_{\mathrm{E}}^2)$) under noisy conditions ($\mathsf{S2CD}$). Table 7.8 presents the restoration results and reviews the performance of our proposed RFI-DRUNet model. These results show that RFI-DRUNet significantly outperforms the random value filling method used in the

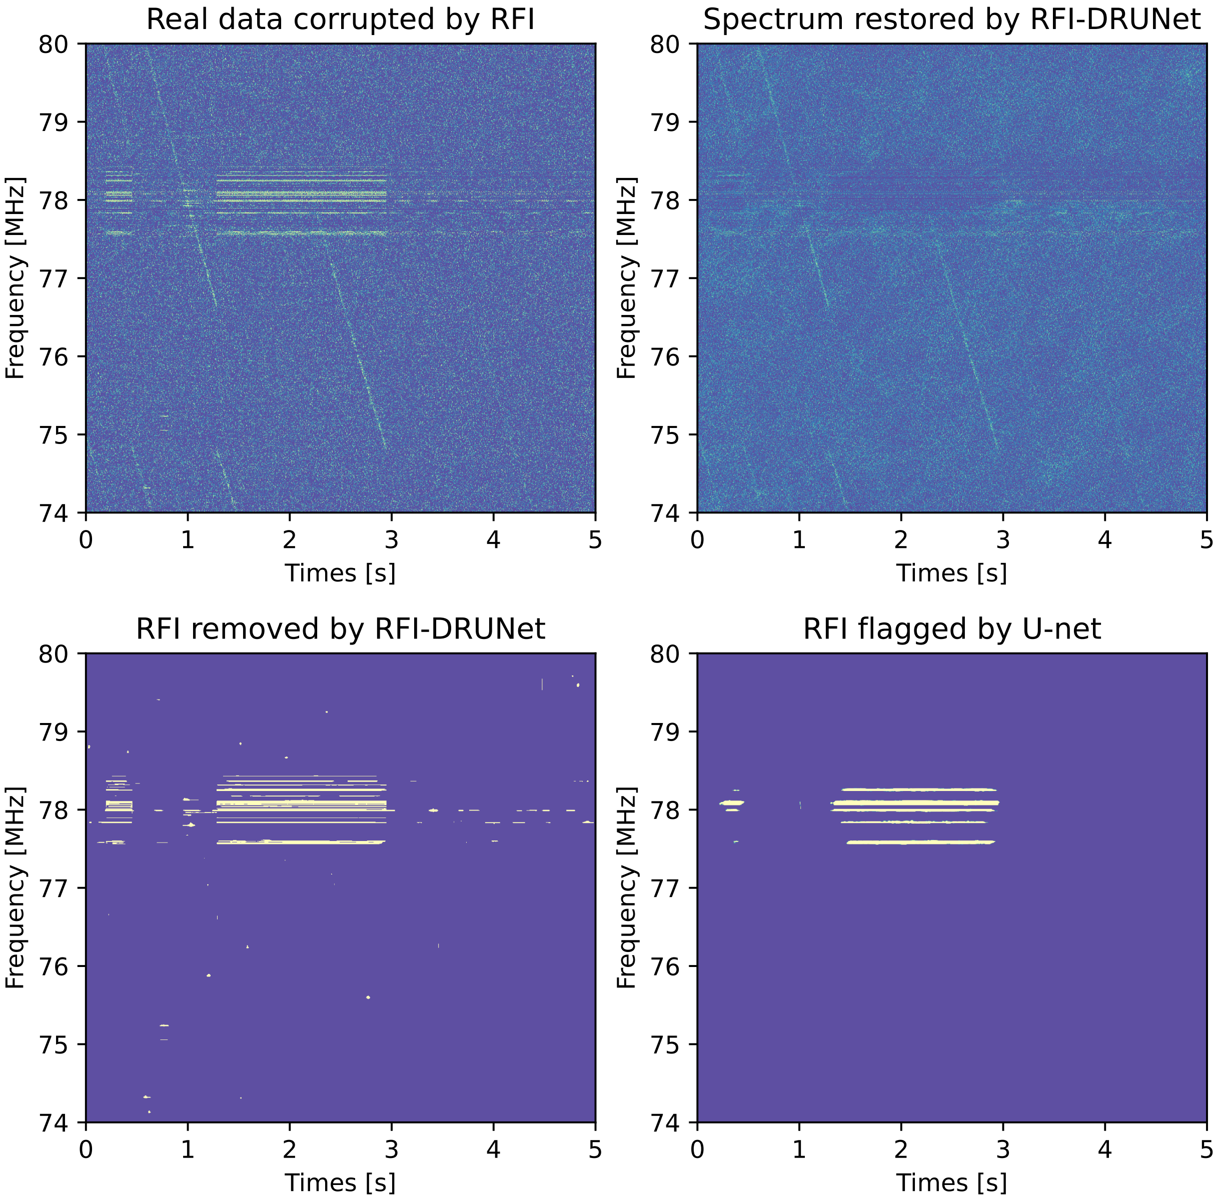

Figure 7.10: Results of visual comparison with RFI-DRUNet and U-Net on a real observation.

comparison strategy in terms of reliability in restoring damaged time-frequency segments.

It is interesting to note that although RFI-DRUNet was originally designed for denoising tasks, where it has excelled, it has also demonstrated comparable performance to U-Net and RFI-Net in RFI detection and marking tasks, as shown in Table 7.6. This ability may be due to the higher demands placed on the model during training for denoising tasks. In fact, RFI-DRUNet network is not more complex than U-Net or RFI-Net, instead it requires fewer network parameters to be learned. The number of parameters of compared models are reported in Table ??. Denoising tasks require the model not only to identify noise, but also to recover clean signals, which typically involves deeper feature learning and more complex network architectures. Consequently, RFI-DRUNet has been designed to handle this complexity, providing it with robust feature extraction and signal restoration capabilities.

Given the similarity between denoising and RFI detection objectives, both of which require the separation of signals of interest from the background, training RFI-DRUNet for denoising tasks may have indirectly improved its performance in RFI detection tasks. In addition, adjusting the training strategies for U-Net and RFI-Net to more closely match the training requirements of the denoising tasks may lead to improved RFI labeling performance. In essence, overly simplistic training strategies may not fully

	Method		**PSNR**
	Detection	**Filling value**	
$\mathrm{S_2C_D}$	Data		31.91 ±1.33
	RFI-DRUNet		59.95 ±4.53
	Oracle	$\mathcal{N}(0, \sigma_{\mathrm{E}}^2)$	42.27 ±1.52
	U-Net	$\mathcal{N}(0, \sigma_{\mathrm{E}}^2)$	41.64 ±1.33
	RFI-Net	$\mathcal{N}(0, \sigma_{\mathrm{E}}^2)$	41.69 ±1.32

Table 7.8: Restoration performance of compared algorithms in terms of average PSNR and standard deviation computed over the test data sets.

	RFI-DRUNet	**RFI-Net**	**U-Net**
♯ parameters ($\times 10^6$)	32.65	48.21	17.26

Table 7.9: Number of parameters of the compared models.

exploit the potential of these models, whereas increasing the complexity of the training tasks may more effectively harness and exploit the capabilities of the models.

Résumé du chapitre

Ce chapitre présente une étude empirique ainsi que des résultats qui mettent en évidence l'efficacité de notre approche RFI-DRUNet dans la suppression des signaux parasites dans les spectres dynamiques. Nous avons mené une série d'expériences approfondies, démontrant que notre méthode est capable de restaurer des spectres de haute qualité, indépendamment de la diversité des intensités et des types de signaux parasites. Notre cadre d'expérience inclut l'entraînement et la validation à l'aide de jeux de données simulés, conçus pour inclure ou exclure le bruit et les différents types de signaux parasites. Huit modèles distincts ont été entraînés, chacun étant spécifiquement conçu pour un scénario de bruit, et une validation croisée a été réalisée pour évaluer leur capacité à récupérer des données affectées par des signaux parasites spécifiques. De plus, des tests ont été effectués sur des ensembles de données avec des niveaux variés de bruit et des données non couvertes par le modèle pour évaluer la résilience des modèles face à des signaux parasites inconnus et à des conditions de bruit changeantes.

Dans la suite de notre expérience, nous avons appliqué les résultats de la restauration à la détection des signaux parasites et avons comparé notre approche à d'autres méthodes de segmentation couramment utilisées pour cette tâche. Les résultats de la comparaison ont révélé que RFI-DRUNet, initialement conçu pour la suppression des signaux parasites, était tout aussi efficace que les méthodes existantes pour la détection de ces signaux.

En outre, nous avons effectué une analyse approfondie de l'architecture du modèle à travers des expériences d'ablation, identifiant les éléments clés qui contribuent le plus à la performance de la restauration. Nous avons également exploré des stratégies pour réduire le nombre de paramètres du réseau, afin d'accélérer le processus d'entraînement et d'inférence tout en réduisant la consommation de mémoire, sans compromettre les performances. L'objectif était d'évaluer la faisabilité du modèle pour des applications pratiques où l'efficacité et l'utilisation des ressources sont primordiales.

Les résultats expérimentaux ont montré que RFI-DRUNet excelle dans la restauration des signaux perturbés par les signaux parasites, dépassent de loin les méthodes de détection de signaux parasites courantes. Notre méthode non seulement identifie avec précision les RFI, mais restaure également de manière fiable les segments de temps-fréquence endommagés. RFI-DRUNet a démontré des performances comparables à celles de U-Net et RFI-Net dans la détection de signaux parasites, ce qui pourrait s'expliquer par les exigences plus élevées imposées au modèle pendant l'entraînement pour la dénoising.

En conclusion, RFI-DRUNet a démontré une performance remarquable dans la réduction des RFI, identifiant et restaurant efficacement les parties du spectre dynamique perturbées.

Chapter 8

Application to pulsar TOA estimation

This chapter will present the application of the proposed approach to pulsar timing. Pulsar timing is a crucial task in pulsar observation that requires extremely high temporal precision. As described in Chapter 1, pulsars have highly stable rotational periods, resulting in the signals from their emissions being received at consistent intervals. The precise measurement of the Times Of Arrival, known as TOAs, is the foundation of pulsar timing, also referred to as chronometry. By analyzing pulsar TOA, researchers can infer a range of intrinsic physical properties of pulsars as well as other astronomical parameters, aiding in studies such as cosmic clocks and the detection of gravitational waves. However, RFI poses a major challenge to the accurate measurement of pulsar signals. By applying our RFI restoration technique, we can effectively mitigate the impact of interference on pulsar signals, thereby significantly enhancing timing accuracy. In the following, we will provide a detailed description of the TOA estimation method and the specific application process and results of our technique in pulsar timing.

Of the many methods used to estimate TOA, the most common and widely used is the Taylor method (Taylor, 1992). The Taylor method accurately determines the arrival time of each pulse by matching the time series of the pulse signal to a predetermined template and employing a least-squares approach to minimize the discrepancy between the observed signal and the template. The method utilizes the stability of the pulsed signal to determine the TOA by incrementally adjusting the time offset of the signal arrivals, aiming to optimize the match. By framing the optimization problem in the Fourier domain, it achieves an estimation accuracy of at least $0.1\delta t$, which is superior to that obtained in the time domain.

In the Taylor method, the pulse signal $p(t)$ is expressed as in relation to a corresponding "standard profile" $s(t)$ through the formulation:

$$p(t) = a + bs(t - \tau) + g(t) \tag{8.1}$$

where a represents an amplitude shift, b is an amplitude scaling factor, and τ is a constant temporal offset. The term $g(t)$ accounts for additional noise, including radiometer and background noise. In addition, the variable t ranges from 0 to P, where P is the topocentric period. The task of determining TOA involves accurately estimating the time shift τ and adding it to the recorded observation start time. In practice, we assume that the observed profile and standard profile $p(t)$ and $s(t)$ are typically sampled at discrete time intervals $t_j = j\Delta t$, $j = 0, 1, \cdots, N-1$ where $\Delta t = P/N$, and N is an integer. Prior to sampling, the detected signals are subjected to a low-pass filter with a cutoff frequency $f_c \leq (2\Delta t)^{-1}$. To ensure that no valuable data is lost during filtering, Δt is selected to be sufficiently small so that f_c exceeds the highest significant frequencies present in the signal. According to the finite sampling theorem, all relevant information is fully and unambiguously captured within the discretely sampled values $p_j = p(t_j)$.

The offset τ can be determined by the method of least squares, the aim being to minimise

$$\chi^2(a, b, \tau) = \sum_{j=0}^{N-1} \left| \frac{P_j - a - bS_{j-\tau}}{\sigma_j} \right|^2 \tag{8.2}$$

where σ_k is the noise term. In this equation, we can easily compute the offset τ, which is an integer multiple of the time resolution Δt. Furthermore, achieving greater precision in delay estimation is more straightforward by transforming the problem into Fourier space. If $p(t)$ is equal to a shift and scaled replica of $s(t)$ plus random noise, as defined in Equation (8.1), their Fourier transforms are also relate in a sample way and can be written as:

$$p_k = P_k \exp(i\theta_k) = R_{p_k} + \jmath I_{p_k} = \sum_{j=0}^{N-1} p_j e^{i2\pi jK/N} \tag{8.3}$$

$$s_k = S_k \exp(i\phi_k) = R_{s_k} + \jmath I_{s_k} = \sum_{j=0}^{N-1} s_j e^{i2\pi jK/N} \tag{8.4}$$

where the frequency index k ranges from 0 to $N-1$. Thus the real quantities P_k and S_k are the amplitude of the complex Fourier coefficients and θ_k and ϕ_k are the phases. By linear transformation, Equation (8.1) can be rewritten as a relationship between P_k and S_k:

$$P_k \exp(i\theta_k) = aN + bS_k \exp[i(\phi_k + k\tau)] + G_k, k = 0, \cdots, (N-1) \tag{8.5}$$

where G_k stands for random noise corresponds to the sampled noise in the time-domain profile $g(t_j)$.

It is important to note that the bias a and the scale factor b play analogous roles in both the time and frequency domains. According to the shift theorem (Bracewell and Kahn, 1966), the temporal offset τ manifests in the frequency domain as a linear ramp, specifically as $k\tau$ added to the phases of the Fourier coefficients of the standard profile. Once the transforms have been computed, the value of a can be directly obtained by:

$$a = (P_0 - bS_0)/N \tag{8.6}$$

The desired pulse time of arrival τ and gain factor b can determined by minimizing the goodness-of-fit statistic

$$\chi^2(b, \tau) = \sum_{k=1}^{N/2} \left| \frac{P_k - bS_k e^{ik\tau}}{\sigma_k} \right|^2 \tag{8.7}$$

In this equation σ_k is the root-mean-square amplitude of the noise at frequency k. Due to the inherent symmetries in the transforms, the limits of the equation can be adjusted to range from 1 to $N/2$ rather than 0 to N

Expressing the Equation (8.7) in complex form and expanding the square modulus yields a simpler expression

$$\chi^2 = \sigma^{-2} \sum \left(P_k P_k^* - bP_k S_k e^{ik\tau} - bP_k S_k^* e^{-ik\tau} + b^2 S_k S_k^* \right) \tag{8.8}$$

The offset τ and amplitude b in the Equation (8.8) can be solved with $\frac{\partial \chi^2}{\partial \tau} = 0$, and $\frac{\partial \chi^2}{\partial b} = 0$

$$f(\tau) = \frac{2b}{\sigma^2} \sum_k \left[(R_{p_k} I_{s_k} - Ip_k R_{s_k}) \cos k\tau + (R_{p_k} R_{s_k} + I_{p_k} I_{s_k}) \sin k\tau \right] = 0 \tag{8.9}$$

$$b = \frac{\sum (R_{p_k} R_{s_k}) \cos k\tau + (I_{p_k} R_{s_k} - R_{s_k} I_{s_k}) \sin k\tau}{\sum R_{s_k}^2 + I_{s_k}^2} \tag{8.10}$$

where Equation (8.10) can be computed directly to yield the value of the amplitude b. The Brent method (Press, 1989) can be employed in Equation (8.9) to find a suitable τ such that $f(\tau) = 0$.

Uncertainties in the estimated values of τ and b can be assessed by approximating χ^2 near its minimum using the leading terms of a Taylor series. By determining the changes in b and τ necessary to increase χ^2 by 1, one can derive the uncertainties. This method results in:

$$\sigma_\tau^2 = \left(\frac{\partial^2 \chi^2}{\partial \tau^2} \right)^{-1} = \frac{\sigma^2}{2b \sum k^2 (R_{P_k} R_{S_k} + I_{P_k} I_{S_k}) \cos k\tau + (I_{P_k} R_{S_k} - R_{P_k} I_{S_k}) \sin k\tau} \tag{8.11}$$

$$\sigma_b^2 = \left(\frac{\partial^2 \chi^2}{\partial b^2} \right)^{-1} = \frac{\sigma^2}{2 \sum S_k^2} = \frac{\sigma^2}{2 \sum R_{S_k}^2 + I_{S_k}^2} \tag{8.12}$$

8.1 Numerical experiments

After introducing the Taylor method, we will show how it can be used to validate the effectiveness of our proposed approach. We will show how the estimation of pulsar TOA can be improved by removing the effects of RFI. To perform this validation, we have chosen to use synthetic signals generated according to the scenario $\mathsf{S}_2\mathsf{C}_\mathrm{D}$ with a pulsar template $\bar{P}(\cdot,\cdot)$ which have been presented in Section 4.3. These simulated data will help us to specify how our RFI repair method can optimise TOA estimation in real applications and to investigate its performance under different interference conditions. The parameters for generating synthetic data are the same as in Section 4.5.2, except that the pulsar period $\rho = 64$ is fixed. According to the assumption of Taylor's method, the observed signal pulsar $P(n,k)$ can be modeled as a noisy, scaled, and temporally shifted version of the template signal denoted $\bar{P}(n,k)$, i.e.,

$$P(n,k) = u + v\bar{P}(n,k-\Delta t) + E(n,k) \tag{8.13}$$

where u is an amplitude offset, b is a scaling factor, Δt is the temporal shift between the observed signal and the template, and $E(k,n)$ is a term accounting for modeling errors.

Taylor's estimation method was applied to compute the TOA values for three different types of signals: the RFI-free signal, the RFI-corrupted signal, and the signal restored by RFI-DRUNet. When generating the pulsar signal, we can obtain the TOA from its pulse profile with the knowledge of the relevant parameters, which we consider to be the true value. By applying Taylor's method to each of these signals, we obtained their respective TOA estimates. To evaluate the accuracy of the TOA estimates, we used two metrics: Mean Squared Error (MSE) and Mean Absolute Error (MAE). These metrics were employed to assess the deviation of the TOA estimates from the true value, thereby demonstrating the accuracy of TOA estimation with and without RFI correction using RFI-DRUNet.

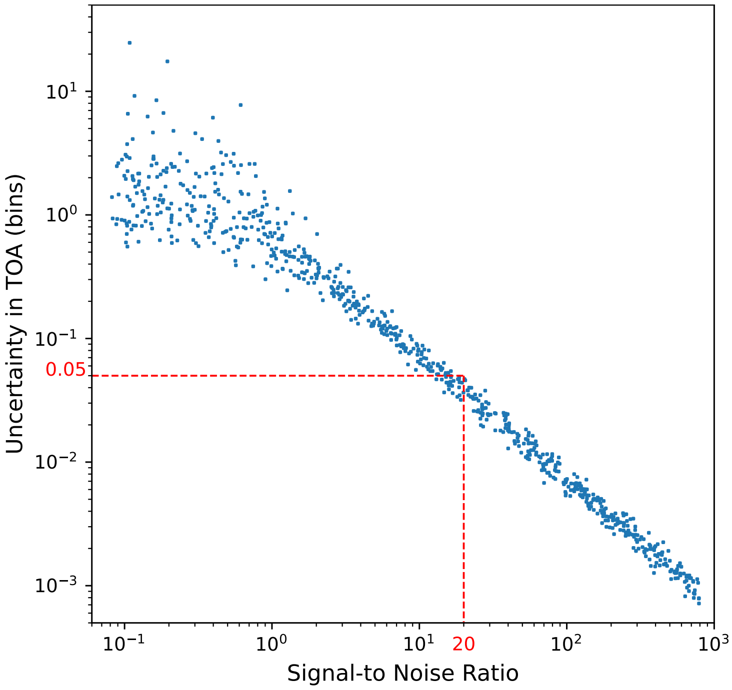

Figure 8.1: Scatter plot of the TOA estimation uncertainty versus SNR. The red dashed lines indicate the maximum SNR value which is 20, considered in Scenario $\mathsf{S}_2\mathsf{C}_\mathrm{D}$,along with the corresponding uncertainty value, adjusted through empirical linear regression in the log-log space.

It is worth noting that Taylor's method provides an uncertainty measure for its TOA estimates, denoted as $\sigma_{\Delta t}$. This uncertainty reflects the difficulty of the estimation task and is influenced by factors such as the SNR: a lower SNR results in higher uncertainty. This additional information provided by the Taylor method can be used to exclude all estimates with uncertainty measures above a given threshold, since in these cases the pulsar signals are hidden by the background noise and therefore these estimates are considered unreliable. Figure. 8.1 provide an illustration of this finding, where each point corresponds to a specific signal generated across a broad range of SNR values (from 10^{-1} to 10^{3}). Each signal is characterized by its SNR (on the x-axis) and the uncertainty measure provided by the estimation method (on the y-axis). In this log-log scatter plot, the uncertainty is evidently inversely proportional to the SNR. According to the parameters used to generate the pulsar signals, the upper limit of the SNR value of the pulsar signals is 20, while the corresponding lower uncertainty limit is 0.05 based on experience.

The Taylor method is applied to RFI-free signals, signals corrupted by RFI and signals recovered by RFI-DRUNet. In addition, the Taylor method will be applied to signals processed by the three traditional RFI detection methods, i.e. RFI signals are located/tagged by Oracle Detector, U-Net or RFI-Net, and points identified as corrupted are replaced by random values. These operation is widely used in current pulsar research to mitigate the impact of RFI and accurately estimate TOA of pulsar signals. To compare the performance of our proposed RFI restoration method with these existing RFI handling techniques in TOA estimation, allowing us to more comprehensively evaluate the effectiveness of our method in improving TOA estimation accuracy. The estimation errors (expressed as MSE and MAE) as a function of the uncertainty threshold are depicted in Figure 8.2. As expected, the error in estimating TOAs from RFI-free dynamic spectra is minimal, as the RFI-free signal contains only the pulsar signals and background noise. In contrast, the TOAs estimated from RFI-corrupted signals show a significant deviation from the theoretical values, which is reasonable given the disruptive impact of RFI on the signal. Most importantly, the TOA values obtained from the signals restored by our proposed method are remarkably close to those derived from the RFI-free signals. This demonstrates that our method effectively mitigates the impact of RFI on TOA estimation. Additionally, the TOA values derived from signals processed by the three traditional RFI detection methods exhibit significantly larger errors compared to those obtained using our proposed method. In some cases, these errors even surpass those from the RFI-corrupted signals. This finding not only demonstrates the superior performance of our approach in TOA estimation but also highlights a major drawback in current methods, where replacing RFI-flagged data with random values or zeros can introduce substantial inaccuracies.

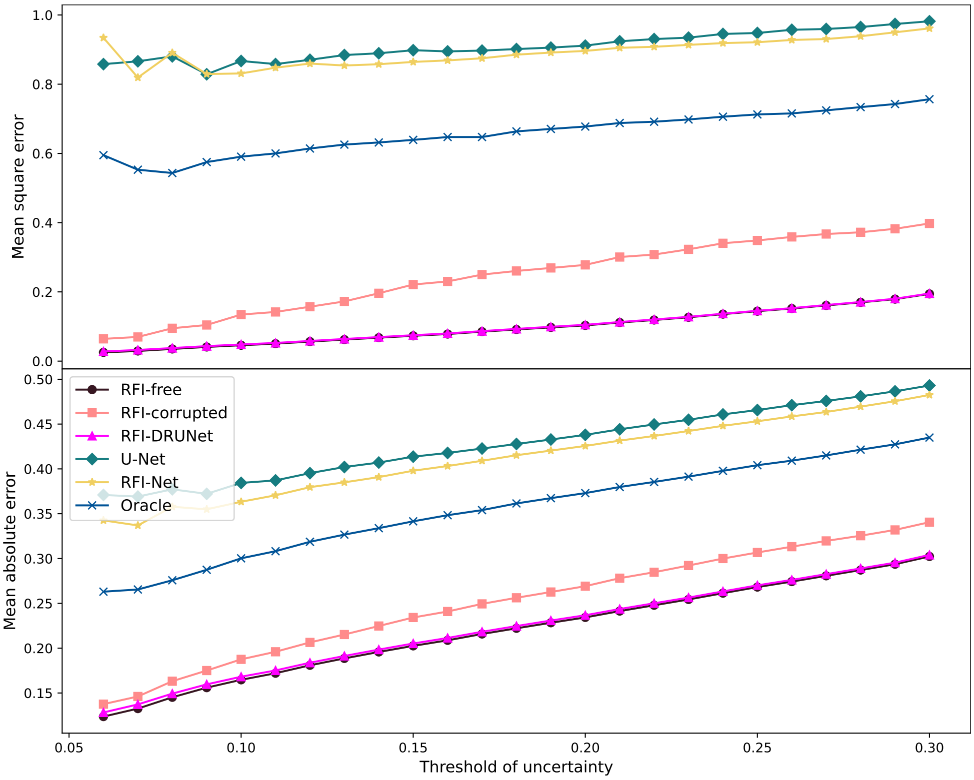

Figure 8.2: TOA estimation errors in terms of MSE (top) and MAR (bottom) as function of the uncertainty threshold obtained from RFI-free (brown dots lines), RFI-corrupted (pink square line), restored by RFI-DRUNet (rose triangle line), restored by U-Net (dark green diamond line), restored by RFI-Net (yellolw star line), and restored by Oracle (dark blue cross line).

Résumé du chapitre

Nous avons développé une approche, intitulée RFI-DRUNet, conçue pour améliorer l'exactitude des estimations de temps d'arrivée (TOA) en identifiant et en réparant les parties du spectre-dynamique corrompues par des RFI. La modélisation des pulsars comme des horloges est une composante essentielle dans l'étude de ces étoiles, nécessitant une précision temporelle extrême. Les pulsars sont caractérisés par des périodes de rotation très stables, ce qui se traduit par des signaux reçus à des intervalles cohérents. La mesure précise de ces intervalles est la base de l'horloge pulsar ou chronométrie. L'analyse des TOA des pulsars permet aux chercheurs d'inférer une variété de propriétés physiques intrinsèques des pulsars, ainsi que d'autres paramètres astronomiques, contribuant ainsi à des recherches telles que les horloges cosmiques et la détection d'ondes gravitationnelles. Cependant, les RFI représentent un défi significatif pour la mesure précise des signaux pulsars. En implémentant notre technique de restauration des RFI, nous pouvons réduire efficacement l'impact des interférences sur les signaux pulsars, améliorant ainsi la précision temporelle.

Pour démontrer l'efficacité de notre approche de réparation des RFI, nous avons utilisé la méthode de Taylor pour estimer les TOA. Cette méthode détermine avec précision l'heure d'arrivée de chaque pulsation en alignant la série temporelle du signal de pulsation avec un modèle préétabli et en utilisant une approche des moindres carrés pour minimiser la divergence entre le signal observé et le modèle. Cette méthode tire parti de la stabilité du signal pulsé pour déterminer le TOA en ajustant progressivement le décalage temporel des arrivées de signaux, dans le but d'optimiser la correspondance. Porter le problème d'optimisation dans le domaine de Fourier permet d'obtenir une précision d'estimation supérieure à celle obtenue dans le domaine temporel.

Nous avons utilisé des signaux synthétiques pour évaluer l'efficacité de notre approche de réparation des RFI, générés selon le scénario S2CD et un modèle de pulsar décrit dans la Section 4.3. Ces données simulées nous aident à déterminer comment notre méthode de réparation des RFI peut optimiser l'estimation du TOA dans des applications réelles et à étudier ses performances sous diverses conditions d'interférence. Les paramètres pour générer des données synthétiques sont identiques à ceux de la section 4.5.2, à l'exception de la période de pulsar ρ qui est fixée à 64 bins.

Nous avons appliqué la méthode de Taylor pour calculer les valeurs de TOA pour trois types de signaux : le signal sans RFI, le signal affecté par les RFI et le signal restauré grâce à RFI-DRUNet. Lors de la génération du signal pulsar, nous pouvons obtenir le TOA pour son profil d'impulsion avec la connaissance des paramètres pertinents, que nous considérons comme la valeur réelle. En appliquant la méthode de Taylor à chacun de ces signaux, nous avons obtenu leurs estimations de TOA respectives. Pour évaluer la précision des estimations de TOA, nous avons utilisé deux métriques : l'erreur quadratique moyenne (MSE) et l'erreur absolue moyenne (MAE). Ces métriques ont été utilisées pour évaluer la déviation des estimations de TOA par rapport à la valeur réelle, démontrant ainsi la précision de l'estimation de TOA avec et sans correction des RFI grâce à RFI-DRUNet.

Il est important de souligner que la méthode de Taylor fournit une mesure d'incertitude pour ses estimations de TOA. Cette incertitude reflète la complexité de la tâche d'estimation et est influencée par des facteurs tels que le rapport signal-sur-bruit (SNR), en particulier un SNR plus faible entraîne une incertitude plus élevée. Cette information supplémentaire fournie par la méthode de Taylor peut être utilisée pour exclure toutes les estimations dont la mesure d'incertitude dépasse un seuil donné, car dans ces cas, les signaux de pulsars sont masqués par le bruit de fond, et par conséquent, ces estimations sont considérées comme non fiables. Nos résultats montrent que RFI-DRUNet dépasse les techniques existantes en matière d'estimation des TOA. Ils montrent en outre l'intérêt de la méthode proposée par rapport aux méthodes de la littérature qui remplacent les données corrompues par des valeurs aléatoires ou nulles, ce qui peut introduire des erreurs d'estimation des TOA substantielles.

Conclusion

The main aim of this work is to apply deep learning techniques to Mitigating Radio Frequency Interference (RFI), especially in the field of pulsar observation. Through the development of innovative algorithms, this work not only provides new perspectives to address this pressing problem, but also enhances the potential for gravitational wave detection by improving the accuracy of pulsar timing, thus contributing significantly to the advancement of radio astronomy.

Our main contribution lies in formulating the RFI mitigation problem into an image recovery task and utilizing deep convolutional networks for the recovery of observations corrupted by RFI. This approach overcomes the limitation of traditional RFI detection and localization methods, namely the inability to reliably recover signals affected by RFI, which usually leaves the option of discarding such affected data. To achieve this goal, we propose a customized network architecture that performs well in image denoising task. In order to train the proposed network, we have designed a simulation framework based on physically inspired and statistical models capable of generating pulsar signals and RFI signals. Using the observational properties of the NenuFAR observatory, we finely set the parameters in the framework to generate the dynamic spectra of pulsar observations corrupted by RFI and their corresponding dynamic spectra unperturbed by RFI. The dataset consisting of these paired dynamical spectra enabled the training of our deep learning denoising network. On these simulated datasets, we conducted extensive experiments and demonstrated that the proposed method not only accurately identifies RFI, but also reliably recovers the disturbed positions. We also tested the generalization ability and robustness of the model, and conducted model analysis and ablation experiments to explore for its deployment. Ultimately, we demonstrate the significant impact of the proposed method on improving timing accuracy in a pulsar timing application.

Our work provides a new solution to the problem of mitigating RFI in radio astronomy, and the validity of our proposed idea and its potential for deployment in practical applications is verified by extensive experiments on simulated datasets. We have used pulsar observations at the NenuFAR observatory as an example, but the idea can be extended to a variety of observations at more observatories. Utilizing various types of networks with superior performance in the field of image restoration tasks can help us to effectively mitigate RFI in various real-world applications, and even develop generic models capable of recovering from RFI damage in various observations.

Directions of future work is in three folds. Firstly, this study mainly focuses on the intensity of the complex signal. It would be relevant to also exploit the information brought by the phase of the signal. However, the main challenge in considering phase information to the training data is the modeling of phase information, especially modeling pulsars in the phase information. Secondly, it can be extended to other observational frequency ranges, as well as to observations of other astronomical objects. Finally, with regard to alternative options for modeling datasets, the use of currently popular generative networks or diffusion models for data generation could be considered. Diffusion models have been shown to be capable of generating high-quality images on natural data, and there has been related work exploring the application of diffusion models for generating images of astronomical observations. However, the complexity of diffusion models is much higher than that of convolutional neural networks, and the required training resources represent a significant challenge.

In this study, we provide a new solution for pulsar observation in radio astronomy by transforming the RFI mitigation into an image restoration and using deep learning techniques. Our work is not only innovative in theory, but also shows remarkable results in practical applications. Through extensive experiments on simulated datasets, we demonstrate the effectiveness of the proposed method and provide a solid foundation for its deployment in practical astronomical observations. In addition, our study opens up the possibility of future applications in a wider range of frequencies and astronomical observations, contributing lasting value to the development of radio astronomy.

Conclusion (français)

Le principal but de cette étude est d'appliquer des techniques d'apprentissage profond à l'atténuation des interférences radioélectriques (RFI), en particulier dans le domaine de l'observation des pulsars. En développant des algorithmes innovants, ce travail offre non seulement de nouvelles perspectives pour résoudre ce problème crucial, mais renforce également le potentiel de détection des ondes gravitationnelles en améliorant la précision du chronométrage des pulsars, contribuant ainsi de manière significative à l'avancement de la radioastronomie.

Notre principale contribution consiste à formuler le problème de l'atténuation des interférences radioélectriques en une tâche de restauration d'images et à utiliser des réseaux convolutionnels profonds pour la correction des observations corrompues par les interférences. Cette approche permet de surmonter les limites des méthodes traditionnelles de détection et de localisation des interférences, à savoir l'incapacité de récupérer de manière fiable les signaux affectés par les interférences, ce qui laisse généralement l'option d'écarter ces données affectées. Pour atteindre cet objectif, nous proposons une architecture de réseau personnalisée qui donne de bons résultats dans les tâches de débruitage d'images. Afin d'entraîner le réseau proposé, nous avons conçu un cadre de simulation basé sur des modèles statistiques guidés par la physique capables de générer des signaux de pulsars et des signaux RFI. En utilisant les propriétés d'observation de l'observatoire NenuFAR, nous avons finement réglé les paramètres pour générer les spectres dynamiques des observations de pulsars corrompus par RFI et leurs spectres dynamiques correspondants non perturbés par RFI. L'ensemble de données constitué de ces spectres dynamiques appariés a permis l'entraînement de notre réseau de débruitage par apprentissage profond. Sur ces ensembles de données simulées, nous avons mené des expériences approfondies et démontré que la méthode proposée non seulement identifie avec précision les interférences, mais estime également de manière fiable les positions perturbées. Nous avons également testé la capacité de généralisation et la robustesse du modèle, et mené des expériences d'analyse et d'ablation de modèle afin d'explorer son déploiement. Enfin, nous démontrons l'impact significatif de la méthode proposée sur l'amélioration de la précision du chronométrage dans une application de chronométrage des pulsars.

Notre travail apporte une nouvelle solution au problème de l'atténuation des interférences en radioastronomie, et la validité de l'idée proposée ainsi que son potentiel de déploiement dans des applications pratiques sont vérifiés par des expériences approfondies sur des ensembles de données simulées. Nous avons pris comme exemple les observations de pulsars à l'observatoire NenuFAR, mais l'idée peut être étendue à une variété d'observations dans d'autres observatoires. L'utilisation de divers types de réseaux aux performances supérieures dans le domaine des tâches de restauration d'images peut nous aider à atténuer efficacement les interférences dans diverses applications réelles, et même à développer des modèles génériques capables de réparer les dommages causés par les interférences dans diverses observations.

Les orientations à considérer pour des travaux futurs sont de trois ordres. Premièrement, cette étude se concentre principalement sur l'intensité du signal complexe. Il conviendrait alors de prendre en compte l'information de phase également. Toutefois, la principale difficulté liée à l'ajout d'informations de phase aux données d'apprentissage est la modélisation des informations de phase, en particulier la modélisation des pulsars dans les informations de phase. Deuxièmement, cette méthode peut être étendue à d'autres gammes de fréquences d'observation, ainsi qu'à des observations d'autres objets astronomiques. Enfin, en ce qui concerne les options alternatives pour la modélisation des ensembles de données, l'utilisation des réseaux génératifs ou des modèles de diffusion actuellement populaires pour la génération de données pourrait être envisagée. Les modèles de diffusion se sont révélés capables de générer des images de haute qualité sur des données naturelles, et des travaux connexes ont exploré l'application des modèles de diffusion à la génération d'images d'observations astronomiques. Cependant, la complexité des modèles de diffusion est beaucoup plus élevée que celle des réseaux neuronaux convolutionnels, et les ressources d'apprentissage requises représentent un défi important.

Dans cette étude, nous proposons une nouvelle solution pour l'observation des pulsars en radioastronomie en transformant l'atténuation des interférences en une restauration d'image et en utilisant des techniques d'apprentissage profond. Notre travail est non seulement innovant sur le plan méthodologique, mais il donne également des résultats remarquables dans les applications pratiques. Grâce à des expériences approfondies sur des ensembles de données simulées, nous démontrons l'efficacité de la méthode proposée et fournissons une base solide pour son déploiement dans des observations astronomiques pratiques. En outre, notre étude ouvre la voie à de futures applications dans une gamme plus large de fréquences et d'observations astronomiques, contribuant ainsi de manière durable au développement de la radioastronomie.

Bibliography

Agazie, G., Anumarlapudi, A., Archibald, A.M., Arzoumanian, Z., Baier, J., Baker, P.T., Bécsy, B., Blecha, L., Brazier, A., Brook, P.R., Burke-Spolaor, S., Burnette, R., Case, R., Casey-Clyde, J.A., Charisi, M., Chatterjee, S., Cohen, T., Cordes, J.M., Cornish, N.J., Crawford, F., Cromartie, H.T., Crowter, K., DeCesar, M.E., DeGan, D., Demorest, P.B., Dolch, T., Drachler, B., Ferrara, E.C., Fiore, W., Fonseca, E., Freedman, G.E., Garver-Daniels, N., Gentile, P.A., Glaser, J., Good, D.C., Gültekin, K., Hazboun, J.S., Jennings, R.J., Johnson, A.D., Jones, M.L., Kaiser, A.R., Kaplan, D.L., Kelley, L.Z., Kerr, M., Key, J.S., Laal, N., Lam, M.T., Lamb, W.G., Lazio, T.J.W., Lewandowska, N., Liu, T., Lorimer, D.R., Luo, J., Lynch, R.S., Ma, C.P., Madison, D.R., McEwen, A., McKee, J.W., McLaughlin, M.A., McMann, N., Meyers, B.W., Mingarelli, C.M.F., Mitridate, A., Natarajan, P., Ng, C., Nice, D.J., Ocker, S.K., Olum, K.D., Pennucci, T.T., Perera, B.B.P., Pol, N.S., Radovan, H.A., Ransom, S.M., Ray, P.S., Romano, J.D., Saffer, A., Sardesai, S.C., Schmiedekamp, A., Schmiedekamp, C., Schmitz, K., Shapiro-Albert, B.J., Siemens, X., Simon, J., Siwek, M.S., Stairs, I.H., Stinebring, D.R., Stovall, K., Sun, J.P., Susobhanan, A., Swiggum, J.K., Taylor, J.A., Taylor, S.R., Turner, E., Unal, C., Vallisneri, M., Vigeland, S.J., Wahl, H.M., Witt, C.A., Young, O., 2023. The NANOGrav 15-year data set: Search for Transverse Polarization Modes in the Gravitational-Wave Background. arXiv e-prints `2310.12138`.

Akeret, J., Chang, C., Lucchi, A., Refregier, A., 2017a. Radio frequency interference mitigation using deep convolutional neural networks. Astronomy and Computing 18, 35–39.

Akeret, J., Seehars, S., Chang, C., Monstein, C., Amara, A., Refregier, A., 2017b. Hide & seek: End-to-end packages to simulate and process radio survey data. Astronomy and Computing 18, 8–17.

An, T., Chen, X., Mohan, P., Lao, B.Q., 2017. Radio frequency interference mitigation. arXiv e-prints `1711.01978`.

Antoniadis, J., Arumugam, P., Arumugam, S., Babak, S., Bagchi, M., Nielsen, A.S.B., Bassa, C.G., Bathula, A., Berthereau, A., Bonetti, M., Bortolas, E., Brook, P.R., Burgay, M., Caballero, R.N., Chalumeau, A., Champion, D.J., Chanlaridis, S., Chen, S., Cognard, I., Dandapat, S., Deb, D., Desai, S., Desvignes, G., Dhanda-Batra, N., Dwivedi, C., Falxa, M., Ferdman, R.D., Franchini, A., Gair, J.R., Goncharov, B., Gopakumar, A., Graikou, E., Grießmeier, J.M., Guillemot, L., Guo, Y.J., Gupta, Y., Hisano, S., Hu, H., Iraci, F., Izquierdo-Villalba, D., Jang, J., Jawor, J., Janssen, G.H., Jessner, A., Joshi, B.C., Kareem, F., Karuppusamy, R., Keane, E.F., Keith, M.J., Kharbanda, D., Kikunaga, T., Kolhe, N., Kramer, M., Krishnakumar, M.A., Lackeos, K., Lee, K.J., Liu, K., Liu, Y., Lyne, A.G., McKee, J.W., Maan, Y., Main, R.A., Mickaliger, M.B., Niţu, I.C., Nobleson, K., Paladi, A.K., Parthasarathy, A., Perera, B.B.P., Perrodin, D., Petiteau, A., Porayko, N.K., Possenti, A., Prabu, T., Leclere, H.Q., Rana, P., Samajdar, A., Sanidas, S.A., Sesana, A., Shaifullah, G., Singha, J., Speri, L., Spiewak, R., Srivastava, A., Stappers, B.W., Surnis, M., Susarla, S.C., Susobhanan, A., Takahashi, K., Tarafdar, P., Theureau, G., Tiburzi, C., van der Wateren, E., Vecchio, A., Krishnan, V.V., Verbiest, J.P.W., Wang, J., Wang, L., Wu, Z., 2023. The second data release from the european pulsar timing array. Astronomy & Astrophysics 678, A50.

Arzoumanian, Z., Baker, P.T., Brazier, A., Burke-Spolaor, S., Chamberlin, S., Chatterjee, S., Christy, B., Cordes, J.M., Cornish, N.J., Crawford, F., et al., 2018. The nanograv 11 year data set: pulsar-timing constraints on the stochastic gravitational-wave background. The Astrophysical Journal 859, 47.

Asad, K., Girard, J., De Villiers, M., Ansah-Narh, T., Iheanetu, K., Smirnov, O., Santos, M., Lehmensiek, R., Jonas, J., De Villiers, D., et al., 2021. Primary beam effects of radio astronomy antennas – II.

modelling MeerKAT L-band beams. Monthly Notices of the Royal Astronomical Society 502, 2970–2983.

Athreya, R., 2009. A new approach to mitigation of radio frequency interference in interferometric data. The Astrophysical Journal 696, 885.

Baade, W., Zwicky, F., 1934. On super-novae. Proceedings of the National Academy of Sciences 20, 254–259.

Baan, W., Fridman, P., Millenaar, R., 2004a. Radio frequency interference mitigation at the westerbork synthesis radio telescope: Algorithms, test observations, and system implementation. The Astronomical Journal 128, 933.

Baan, W., Fridman, P., Millenaar, R., 2004b. Radio frequency interference mitigation with real-time fpga digital signal processing, in: Proc. European Signal Processing Conference (EUSIPCO), IEEE. pp. 765–768.

Baan, W.A., 2011. Rfi mitigation in radio astronomy, in: 2011 XXXth URSI General Assembly and Scientific Symposium, IEEE. pp. 1–2.

Backer, D.C., Kulkarni, S.R., Heiles, C., Davis, M., Goss, W., 1982. A millisecond pulsar. Nature 300, 615–618.

Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. IEEE transactions on pattern analysis and machine intelligence 39, 2481–2495.

Barnbaum, C., Bradley, R.F., 1998. A new approach to interference excision in radio astronomy: Real-time adaptive cancellation. The Astronomical Journal 116, 2598.

Bhat, N., Cordes, J., Chatterjee, S., Lazio, T., 2005. Radio frequency interference identification and mitigation using simultaneous dual-station observations. Radio Science 40, 1–16.

Bondonneau, L., Grießmeier, J.M., Theureau, G., Cognard, I., Brionne, M., Kondratiev, V., Bilous, A., McKee, J., Zarka, P., Viou, C., et al., 2021. Pulsars with nenufar: Backend and pipelines. Astronomy & Astrophysics 652, A34.

Bracewell, R., Kahn, P.B., 1966. The fourier transform and its applications. American Journal of Physics 34, 712–712.

Briggs, F., Bell, J., Kesteven, M., 2000. Removing radio interference from contaminated astronomical spectra using an independent reference signal and closure relations. The Astronomical Journal 120, 3351.

Brossard, M., El Korso, M.N., Pesavento, M., Boyer, R., Larzabal, P., Wijnholds, S.J., 2018. Parallel multi-wavelength calibration algorithm for radio astronomical arrays. Signal Processing 145, 258–271.

Buist, P.J., Engelen, S., Noroozi, A., Sundaramoorthy, P., Verhagen, S., Verhoeven, C., 2011. Overview of pulsar navigation: Past, present and future trends. Navigation 58, 153–164.

Burgay, M., D'Amico, N., Possenti, A., Manchester, R., Lyne, A., Joshi, B., McLaughlin, M., Kramer, M., Sarkissian, J., Camilo, F., et al., 2003. An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system. Nature 426, 531–533.

Chadwick, J., 1932. Possible existence of a neutron. Nature 129, 312–312.

Chang, A., Knapp, M., LaBelle, J., Swoboda, J., Volz, R., Erickson, P.J., 2023. Removing radio frequency interference from auroral kilometric radiation with stacked autoencoders, in: Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (ICASSP), IEEE. pp. 1–5.

Chen, L.C., 2017. Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 .

Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L., 2014. Semantic image segmentation with deep convolutional nets and fully connected crfs. arXiv preprint arXiv:1412.7062 .

Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L., 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence 40, 834–848.

Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H., 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proceedings of the European conference on computer vision (ECCV), pp. 801–818.

Davis, M., Taylor, J., Weisberg, J., Backer, D., 1985. High-precision timing observations of the millisecond pulsar psr1937+ 21. Nature 315, 547–550.

DeBoer, D.R., Parsons, A.R., Aguirre, J.E., Alexander, P., Ali, Z.S., Beardsley, A.P., Bernardi, G., Bowman, J.D., Bradley, R.F., Carilli, C.L., et al., 2017. Hydrogen epoch of reionization array (HERA). Publications of the Astronomical Society of the Pacific 129, 045001.

Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, Ieee. pp. 248–255.

Dewdney, P.E., Hall, P.J., Schilizzi, R.T., Lazio, T.J.L., 2009. The square kilometre array. Proceedings of the IEEE 97, 1482–1496.

Ellingson, S.W., 2005. Rfi mitigation and the ska. The Square Kilometre Array: An Engineering Perspective , 261–267.

Ellingson, S.W., Bunton, J.D., Bell, J.F., 2000. Cancellation of glonass signals from radio astronomy data, in: Radio Telescopes, SPIE. pp. 400–407.

Everett, J., Weisberg, J., 2001. Emission beam geometry of selected pulsars derived from average pulse polarization data. The Astrophysical Journal 553, 341.

Finlay, C., Bassett, B.A., Kunz, M., Oozeer, N., 2023. Trajectory based RFI subtraction and calibration for radio interferometry. arXiv e-prints `2301.04188`.

Fridman, P., 2008. Statistically stable estimates of variance in radio-astronomy observations as tools for radio-frequency interference mitigation. The Astronomical Journal 135, 1810.

Fridman, P., Baan, W., 2001. Rfi mitigation methods in radio astronomy. Astronomy & Astrophysics 378, 327–344.

Fukushima, K., 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. Biological cybernetics 36, 193–202.

Ghanney, Y., Ajib, W., 2020. Radio frequency interference detection using deep learning, in: Proc. IEEE Vehicular Technology Conference (VTC-Spring), IEEE. pp. 1–5.

Gil, J., Gronkowski, P., Rudnicki, W., 1984. Geometry of the emission region of psr 0950+ 08. Astronomy and Astrophysics (ISSN 0004-6361), vol. 132, no. 2, March 1984, p. 312-316. 132, 312–316.

Gold, T., 1968. Rotating neutron stars as the origin of the pulsating radio sources. Nature 218, 731–732.

Goldreich, P., Julian, W.H., 1969. Pulsar electrodynamics. Astrophysical Journal, vol. 157, p. 869 157, 869.

Gu, F., Hao, L., Liang, B., Feng, S., Wei, S., Dai, W., Xu, Y., Li, Z., Dao, Y., 2024. Radio frequency interference detection using efficient multiscale convolutional attention UNet. Monthly Notices of the Royal Astronomical Society 529, 4719–4727. `https://academic.oup.com/mnras/article-pdf/529/4/4719/57147018/stae868.pdf`.

Gullahorn, G.E., Rankin, J.M., 1978. Pulsar timing results from arecibo observatory. Astronomical Journal, vol. 83, Oct. 1978, p. 1219-1224. 83, 1219–1224.

Gunn, J.E., Ostriker, J.P., 1969. Magnetic dipole radiation from pulsars. Nature 221, 454–456.

van Haarlem, M.P., Wise, M.W., Gunst, A., Heald, G., McKean, J.P., Hessels, J.W., de Bruyn, A.G., Nijboer, R., Swinbank, J., Fallows, R., et al., 2013. Lofar: The low-frequency array. Astronomy & astrophysics 556, A2.

Hamid, A., van Straten, W., Griffin, A., 2022. PSRFINET: Radio frequency interference detection in pulsar data with deep residual networks, in: Proc. RFI Workshop.

Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H., 2009. The elements of statistical learning: data mining, inference, and prediction. volume 2. Springer.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.

Hessels, J.W., Ransom, S.M., Stairs, I.H., Freire, P.C., Kaspi, V.M., Camilo, F., 2006. A radio pulsar spinning at 716 hz. Science 311, 1901–1904.

Hewish, A., Bell, S.J., Pilkington, J.D.H., Scott, P.F., Collins, R.A., 1968. Observation of a Rapidly Pulsating Radio Source. Nature 217, 709–713.

Hewish, A., Okoye, S., 1965. Evidence for an unusual source of high radio brightness temperature in the crab nebula. Nature 207, 59–60.

Hinton, G.E., 2012. Neural networks for machine learning. Coursera online course. Lecture 6a - Introduction to Stochastic Gradient Descent.

Hobbs, G., Coles, W., Manchester, R., Keith, M., Shannon, R.M., Chen, D., Bailes, M., Bhat, N., Burke-Spolaor, S., Champion, D., et al., 2012. Development of a pulsar-based time-scale. Monthly Notices of the Royal Astronomical Society 427, 2780–2787.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4700–4708.

Hulse, R.A., Taylor, J.H., 1975. Discovery of a pulsar in a binary system. Astrophysical Journal, vol. 195, Jan. 15, 1975, pt. 2, p. L51-L53. 195, L51–L53.

Ioffe, S., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 .

Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A., 2017. Image-to-image translation with conditional adversarial networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1125–1134.

Jansky, K.G., 1933. Electrical disturbances apparently of extraterrestrial origin. Proceedings of the Institute of Radio Engineers 21, 1387–1398.

Kaspi, V., Helfand, D., 2002. Constraining the birth events of neutron stars, in: ASP Conf. Ser., p. 3.

Kerrigan, J., Plante, P.L., Kohn, S., Pober, J.C., Aguirre, J., Abdurashidova, Z., Alexander, P., Ali, Z.S., Balfour, Y., Beardsley, A.P., et al., 2019. Optimizing sparse RFI prediction using deep learning. Monthly Notices of the Royal Astronomical Society 488, 2605–2615.

Kesteven, M., Hobbs, G., Clement, R., Dawson, B., Manchester, R., Uppal, T., 2005. Adaptive filters revisited: Radio frequency interference mitigation in pulsar observations. Radio Science 40.

Kingma, D.P., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .

Kocz, J., Briggs, F., Reynolds, J., 2010. Radio frequency interference removal through the application of spatial filtering techniques on the Parkes multibeam receiver. The Astronomical Journal 140, 2086.

Kramer, M., Stairs, I.H., Manchester, R., McLaughlin, M., Lyne, A., Ferdman, R., Burgay, M., Lorimer, D., Possenti, A., D'Amico, N., et al., 2006. Tests of general relativity from timing the double pulsar. science 314, 97–102.

Kramer, M., Stairs, I.H., Manchester, R.N., Wex, N., Deller, A.T., Coles, W.A., Ali, M., Burgay, M., Camilo, F., Cognard, I., et al., 2021. Strong-field gravity tests with the double pulsar. Physical Review X 11, 041050.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems 25.

Lazarus, P., Karuppusamy, R., Graikou, E., Caballero, R., Champion, D., Lee, K., Verbiest, J., Kramer, M., 2016. Prospects for high-precision pulsar timing with the new Effelsberg PSRIX backend. Monthly Notices of the Royal Astronomical Society 458, 868–880.

LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE 86, 2278–2324.

Lefkimmiatis, S., 2017. Non-local color image denoising with convolutional neural networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3587–3596.

Li, S.Z., 2009. Markov random field modeling in image analysis. Springer Science & Business Media.

Li, Z., Yu, C., Xiao, J., Long, M., Cui, C., 2021. Detection of radio frequency interference using an improved generative adversarial network. Astronomy and Computing 36, 100482.

Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3431–3440.

Lorimer, D.R., 2008. Binary and millisecond pulsars. Living reviews in relativity 11, 1–90.

Lorimer, D.R., Kramer, M., 2004. Handbook of Pulsar Astronomy. volume 4.

Ma, P.X., Croft, S., Lintott, C., Siemion, A.P.V., 2023. A deep neural network based reverse radio spectrogram search algorithm. RAS Techniques and Instruments 3, 33–43. `https://academic.oup.com/rasti/article-pdf/3/1/33/55568195/rzad056.pdf`.

Manchester, R.N., Hobbs, G.B., Teoh, A., Hobbs, M., 2005. The australia telescope national facility pulsar catalogue. The Astronomical Journal 129, 1993.

Maslakovic, S., Linscott, I., Oslick, M., Twicken, J., 1996. Excising radio frequency interference using the discrete wavelet transform, in: Proc. IEEE Int. Symp. Time-Frequency Time-Scale Analysis (TFTS), IEEE. pp. 349–352.

Matsakis, D.N., Taylor, J.H., Eubanks, T.M., 1997. A statistic for describing pulsar and clock stabilities. Astronomy and Astrophysics, v. 326, p. 924-928 326, 924–928.

Meltzer, D.W., Thorne, K.S., 1966. Normal modes of radial pulsation of stars at the end point of thermonuclear evolution. Astrophysical Journal, vol. 145, p. 514 145, 514.

Mesarcik, M., Boonstra, A.J., Meijer, C., Jansen, W., Ranguelova, E., van Nieuwpoort, R.V., 2020. Deep learning assisted data inspection for radio astronomy. Monthly Notices of the Royal Astronomical Society 496, 1517–1529.

Mesarcik, M., Boonstra, A.J., Ranguelova, E., van Nieuwpoort, R.V., 2022. Learning to detect radio frequency interference in radio astronomy without seeing it. Monthly Notices of the Royal Astronomical Society 516, 5367–5378.

Metz, C.E., 1978. Basic principles of roc analysis, in: Seminars in nuclear medicine, Elsevier. pp. 283–298.

Mhiri, Y., El Korso, M.N., Breloy, A., Larzabal, P., 2022. Multifrequency array calibration in presence of radio frequency interferences. Signal Processing 199, 108613.

Mhiri, Y., El Korso, M.N., Breloy, A., Larzabal, P., 2024. Regularized maximum likelihood estimation for radio interferometric imaging in the presence of radiofrequency interferences. Signal Processing 220, 109430.

Mignani, R.P., Paladino, R., Rudak, B., Zajczyk, A., Corongiu, A., de Luca, A., Hummel, W., Possenti, A., Geppert, U., Burgay, M., et al., 2017. The first detection of a pulsar with alma. The Astrophysical Journal Letters 851, L10.

Milanfar, P., Delbracio, M., 2024. Denoising: A powerful building-block for imaging, inverse problems, and machine learning. arXiv e-prints `2409.06219`.

Morello, V., Barr, E., Cooper, S., Bailes, M., Bates, S., Bhat, N., Burgay, M., Burke-Spolaor, S., Cameron, A., Champion, D., et al., 2019. The high time resolution universe survey – XIV. Discovery of 23 pulsars through GPU-accelerated reprocessing. Monthly Notices of the Royal Astronomical Society 483, 3673–3685.

Nesterov, Y., 1983. A method for solving the convex programming problem with convergence rate o (1/k2), in: Dokl akad nauk Sssr, p. 543.

Noh, H., Hong, S., Han, B., 2015. Learning deconvolution network for semantic segmentation, in: Proceedings of the IEEE international conference on computer vision, pp. 1520–1528.

Offringa, A., De Bruyn, A., Biehl, M., Zaroubi, S., Bernardi, G., Pandey, V., 2010a. Post-correlation radio frequency interference classification methods. Monthly Notices of the Royal Astronomical Society 405, 155–167.

Offringa, A., Van De Gronde, J., Roerdink, J., 2012. A morphological algorithm for improving radio-frequency interference detection. Astronomy & Astrophysics 539, A95.

Offringa, A.R., de Bruyn, A.G., Zaroubi, S., Biehl, M., 2010b. A LOFAR RFI detection pipeline and its first results, in proceedings of rfi mitigation workshop, in: RFI mitigation workshop - RFI2010.

O.M., S., R., S., 2023. Detection and segmentation of radio frequency interference from satellite images using attention-gans. Astronomy and Computing 45, 100769.

Oppenheimer, J.R., Volkoff, G.M., 1939. On massive neutron cores. Physical Review 55, 374.

Ostriker, J., Gunn, J., 1969. On the nature of pulsars. i. theory. Astrophysical Journal, vol. 157, p. 1395 157, 1395.

Owfi, A., Afghah, F., 2023. Autoencoder-based radio frequency interference mitigation for smap passive radiometer, in: IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium, IEEE. pp. 6783–6786.

Pacini, F., 1967. Energy emission from a neutron star. Nature 216, 567–568.

Page, E.S., 1954. Continuous inspection schemes. Biometrika 41, 100–115.

Peck, L.W., Fenech, D.M., 2013. SERPent: Automated reduction and rfi-mitigation software for e-merlin. Astronomy and Computing 2, 54–66.

Pen, U.L., Chang, T.C., Hirata, C.M., Peterson, J.B., Roy, J., Gupta, Y., Odegova, J., Sigurdson, K., 2009. The GMRT EoR experiment: limits on polarized sky brightness at 150 MHz. Monthly Notices of the Royal Astronomical Society 399, 181–194.

Pilkington, J., Hewish, A., Bell, S., Cole, T., 1968. Observations of some further pulsed radio sources. Nature 218, 126–129.

Pitkin, M., 2018. psrqpy: a python interface for querying the ATNF pulsar catalogue. Journal of Open Source Software 3, 538.

Press, W.H., 1989. Numerical recipes in Pascal: the art of scientific computing. volume 1. Cambridge university press.

RA, R.I.R., 2003. Protection criteria used for radio astronomical measurements .

Radhakrishnan, V., Cooke, D., 1969. Magnetic poles and the polarization structure of pulsar radiation .

Reardon, D.J., Shannon, R.M., Cameron, A.D., Goncharov, B., Hobbs, G., Middleton, H., Shamohammadi, M., Thyagarajan, N., Bailes, M., Bhat, N., et al., 2021. The parkes pulsar timing array second data release: timing analysis. Monthly Notices of the Royal Astronomical Society 507, 2137–2153.

Reardon, D.J., Zic, A., Shannon, R.M., Hobbs, G.B., Bailes, M., Marco, V.D., Kapur, A., Rogers, A.F., Thrane, E., Askew, J., Bhat, N.D.R., Cameron, A., Curyło, M., Coles, W.A., Dai, S., Goncharov, B., Kerr, M., Kulkarni, A., Levin, Y., Lower, M.E., Manchester, R.N., Mandow, R., Miles, M.T., Nathan, R.S., Osłowski, S., Russell, C.J., Spiewak, R., Zhang, S., Zhu, X.J., 2023. Search for an isotropic gravitational-wave background with the parkes pulsar timing array. The Astrophysical Journal Letters 951, L6.

Richards, D.W., Comella, J., 1969. The period of pulsar np 0532. Nature 222, 551–552.

Robbins, H., Monro, S., 1951. A stochastic approximation method. The annals of mathematical statistics , 400–407.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18, Springer. pp. 234–241.

Rosenblatt, F., 1958. The perceptron: a probabilistic model for information storage and organization in the brain. Psychological review 65, 386.

Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning representations by back-propagating errors. nature 323, 533–536.

Rutherford, E., 1911. Lxxix. the scattering of α and β particles by matter and the structure of the atom. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 21, 669–688.

Ryle, M., Hewish, A., 1960. The synthesis of large radio telescopes. Monthly Notices of the Royal Astronomical Society 120, 220–230.

Saliwanchik, B.R.B., Slosar, A., 2022. A self-learning neural network approach for radio frequency interference detection and removal in radio astronomy. Publications of the Astronomical Society of the Pacific 134, 114503.

Sclocco, A., Vohl, D., van Nieuwpoort, R.V., 2019. Real-time RFI mitigation for the Apertif Radio Transient System, in: 2019 RFI Workshop-Coexisting with Radio Frequency Interference (RFI), IEEE. pp. 1–8.

Shan, H., 2023. Robust RFI excision for pulsar signals by a novel nonlinear m-type estimator with an application to pulsar timing. The Astrophysical Journal 952, 70.

Shuai, P., 2021. X-ray Pulsar-Based Navigation: Theories and Experiments. Springer Singapore, Singapore. pp. 273–410.

Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 .

Song, Y., Liu, Z., Wang, N., Li, J., Yuen, R., 2021. An algorithm for mitigating transient rfi in pulsar observation. The Astrophysical Journal 922, 94.

Spitkovsky, A., Arons, J., 2004. Time dependence in relativistic collisionless shocks: theory of the variable "wisps" in the crab nebula. The Astrophysical Journal 603, 669.

Staelin, D.H., Reifenstein III, E.C., 1968. Pulsating radio sources near the crab nebula. Science 162, 1481–1483.

Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning, in: Proceedings of the AAAI conference on artificial intelligence.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1–9.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2818–2826.

Tan, M., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946 .

Taylor, J.H., 1991. Millisecond pulsars: Nature's most stable clocks. Proceedings of the IEEE 79, 1054–1062.

Taylor, J.H., 1992. Pulsar timing and relativistic gravity. Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences 341, 117–134.

Taylor, J.H., Fowler, L.A., McCulloch, P.M., 1979. Measurements of general relativistic effects in the binary pulsar psr1913+ 16. Nature 277, 437–440.

Taylor, J.H., Weisberg, J.M., 1982. A new test of general relativity-gravitational radiation and the binary pulsar psr 1913+ 16. Astrophysical Journal, Part 1, vol. 253, Feb. 15, 1982, p. 908-920. NSF-supported research. 253, 908–920.

Vafaei Sadr, A., Bassett, B.A., Oozeer, N., Fantaye, Y., Finlay, C., 2020. Deep learning improves identification of radio frequency interference. Monthly Notices of the Royal Astronomical Society 499, 379–390.

Vila, S.C., 1969. Nature of pulsars. Nature 224, 157–159.

Vos, E.E., Luus, P.F., Finlay, C.J., Bassett, B.A., 2019. A generative machine learning approach to RFI mitigation for radio astronomy, in: Proc. IEEE Int. Workshop Machine Learning Signal Processing (MLSP), IEEE. pp. 1–6.

Wang, H.F., Yuan, M., Yin, Q., Guo, P., Zhu, W.W., Li, D., Feng, S.B., 2020. Radio frequency interference mitigation using pseudoinverse learning autoencoders. Research in Astronomy and Astrophysics 20, 114.

Wang, J., Bacharach, L., Larzabal, P., El Korso, M.N., . A comparison of antenna placement criteria based on the cramér-rao and barankin bounds for interferometer arrays. Available at SSRN 4549498 .

Weisberg, J.M., Nice, D.J., Taylor, J.H., 2010. Timing measurements of the relativistic binary pulsar psr b1913+ 16. The Astrophysical Journal 722, 1030.

Weisberg, J.M., Taylor, J.H., 2004. Relativistic binary pulsar b1913+ 16: Thirty years of observations and analysis. arXiv preprint astro-ph/0407149 .

Werbos, P., 1974. Beyond regression: New tools for prediction and analysis in the behavioral sciences. PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA .

Winkel, B., Kerp, J., Stanko, S., 2007. RFI detection by automated feature extraction and statistical analysis. Astronomische Nachrichten: Astronomical Notes 328, 68–79.

Wootten, A., Thompson, A.R., 2009. The atacama large millimeter/submillimeter array. Proceedings of the IEEE 97, 1463–1471.

Yan, R.Q., Dai, C., Liu, W., Li, J.X., Chen, S.Y., Yu, X.C., Zuo, S.F., Chen, X.L., 2021. Radio frequency interference detection based on the AC-UNet model. Research in Astronomy and Astrophysics 21, 119.

Yang, Z., Yu, C., Xiao, J., Zhang, B., 2020. Deep residual detection of radio frequency interference for FAST. Monthly Notices of the Royal Astronomical Society 492, 1421–1431.

Zarka, P., Coffre, A., Denis, L., Dumez-Viou, C., Girard, J., Grießmeier, J.M., Loh, A., Tagger, M., 2018. The low-frequency radiotelescope nenufar, in: 2018 2nd URSI Atlantic Radio Science Meeting (AT-RASC), pp. 1–1.

Zhang, A., Lipton, Z.C., Li, M., Smola, A.J., 2023. Dive into Deep Learning. Cambridge University Press. `https://D2L.ai`.

Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., Timofte, R., 2021. Plug-and-play image restoration with deep denoiser prior. IEEE Transactions on Pattern Analysis and Machine Intelligence 44, 6360–6376.

Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L., 2017. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Transactions on Image Processing 26, 3142–3155.

Zhang, K., Zuo, W., Zhang, L., 2018. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. IEEE Transactions on Image Processing 27, 4608–4622.

Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y., 2020. Residual dense network for image restoration. IEEE Transactions on Pattern Analysis and Machine Intelligence 43, 2480–2495.

Zhao, J., Zou, X., Weng, F., 2013. WindSat radio-frequency interference signature and its identification over Greenland and Antarctic. IEEE Transactions on Geoscience and Remote Sensing 51, 4830–4839.

Xiao ZHANG

Restauration de spectres dynamiques corrompus par des interférences radioélectriques - Application à l'observation de pulsars par NenuFAR, précurseur basse-fréquence de SKA

Résumé :
Les pulsars sont des étoiles à neutrons qui tournent rapidement, émettent des ondes électromagnétiques et dont nous recevons des impulsions périodiques, agissant ainsi comme des phares cosmiques. L'extrême stabilité de ces signaux périodiques en font des outils inestimables, par exemple pour tester la théorie de la relativité générale ou détecter les ondes gravitationnelles basse fréquence. Cependant, leur observation en radio peut être rendue difficile à cause d'interférence radioélectrique (RFI) d'origine diverse. Ces interférences peuvent masquer les signaux des pulsars, ce qui complique significativement la détection et l'extraction précise de leur signature. Cette thèse vise à proposer des techniques d'apprentissage profond pour atténuer l'impact des RFI dans les observations de pulsars. En particulier, cette thèse illustre l'intérêt de ces méthodes dans le cas d'observations menées par NenuFAR, un précurseur pour le segment basse fréquence du futur Square Kilometre Array (SKA).

La première partie de cette thèse commence par présenter le contexte astrophysique des pulsars, insistant sur leur importance et leur rôle dans l'astronomie contemporaine. Elle présente ensuite le radiotélescope NenuFAR, en détaillant ses spécifications techniques, ses capacités d'observation et ses avantages uniques dans les observations des pulsars à basse fréquence. Enfin, elle passe en revue et résume les méthodes d'atténuation des RFI en radioastronomie, en analysant les limitations des technologies existantes et en identifiant les principaux problèmes auxquels cette thèse vise à répondre. En particulier, les méthodes actuelles de traitement des RFI conduisent souvent à la suppression des données identifiées comme affectées par des RFI, ce qui entraîne inévitablement une perte d'informations précieuses.

Dans la deuxième partie de cette thèse, nous présentons d'abord un cadre de simulation conçu pour générer des données d'observation de pulsars corrompues par des RFI. L'objectif principal de ce cadre est de créer un ensemble de données qui permettant l'entraînement des modèles d'apprentissage profond, contournant alors le manque récurrent de données réelles correctement étiquetées.

Après une revue approfondie de l'apprentissage profond et des réseaux de neurones convolutifs, nous soulignons les insuffisances des méthodes de traitement des RFI existantes, qui impliquent généralement l'identification et la suppression des données suspectées d'être perturbées par les RFI. Pour éviter cette perte d'information relever ce défi, nous proposons une nouvelle approche qui consiste à formuler le problème d'atténuation des RFI comme une tâche de restauration d'image, visant à reconstruire les signaux altérés par les RFI. En capitalisant sur des avancées récentes en apprentissage profond, nous montrons que cette tâche peut être réalisée efficacement grâce à un réseau, baptisé RFI-DRUnet, dont l'architecture est inspirée d'un réseau débruiteur performant.

Dans les chapitres suivants, nous illustrons l'applicabilité et l'efficacité de l'approche proposée sur des données d'observation de pulsars simulées et réelles, la comparant aux approches actuelles basées sur les réseaux neuronaux convolutifs profonds qui détectent les RFI. Les résultats montrent que la méthode proposée peut restaurer efficacement les signaux d'observation altérés par les RFI sous diverses conditions. En outre, nous discutons l'impact potentiel de cette méthode pour améliorer la chronométrie des pulsars.

Restoring dynamic spectra corrupted by radio frequency interference - Application to pulsar observation by SKA-pathfinder low-frequency NenuFAR

Abstract :
Pulsars are rapidly rotating neutron stars that emit electromagnetic radiation signals received as periodic pulses, thereby acting as cosmic lighthouses. The excellent stability of these periodic signals makes them invaluable tools for assessing the theory of general relativity and for detecting low frequency gravitational waves. However, their observation by radio telescopes can be made difficult by radio frequency interference (RFI) of various origins. These interferences can alter the pulsar signals, which may significantly complicate the detection and precise extraction of their signatures. This thesis aims to propose deep learning techniques to mitigate the impact of RFI in pulsar observations. In particular, this thesis illustrates the interest of these methods in the case of observations carried out by NenuFAR, a pathfinder for the low-frequency segment of the upcoming Square Kilometre Array (SKA). The first part of this thesis begins presenting the astrophysical background of pulsars, emphasizing their importance and role in contemporary astronomy. It then introduces the NenuFAR radio telescope, detailing its technical specifications, observational capabilities, and unique advantages in low-frequency pulsar observations. Lastly, it reviews and summarizes RFI mitigation methods in radio astronomy, analyzing the limitations of existing technologies and identifying the key issues this thesis aims to address. In particular, current RFI mitigation methods often lead to the deletion of data identified as affected by RFI, which inevitably results in a loss of valuable information. In the second part of this thesis, we first present a simulation framework designed to generate RFI-corrupted pulsar observation data. The main goal of this framework is to build a dataset that allows the training of deep learning models, thus circumventing the recurring lack of properly labeled real data. After a thorough review of deep learning and convolutional neural networks, we highlight the shortcomings of existing RFI processing methods, which typically involve identifying and removing data suspected of being corrupted by RFI. To avoid this information loss and address this challenge, we propose a novel approach that consists in formulating the RFI attenuation problem as an image restoration task, aiming at reconstructing RFI-altered signals. Capitalizing on recent advances in deep learning, we show that this task can be efficiently achieved using a network, named RFI-DRUnet, whose architecture is inspired by a high-performance denoising network. In the following chapters, we illustrate the applicability and effectiveness of the proposed approach on simulated and real pulsar observation data, comparing it to current approaches based on deep convolutional neural networks that detect RFI. The results show that the proposed method can effectively restore RFI-corrupted observational signals under various conditions. Furthermore, we discuss the potential impact of this method for improving pulsar chronometry.

Laboratory of Physics and Chemistry of the Environment
and Space (LPC2E)
3 Av. de la Recherche Scientifique, 45071 Orléans