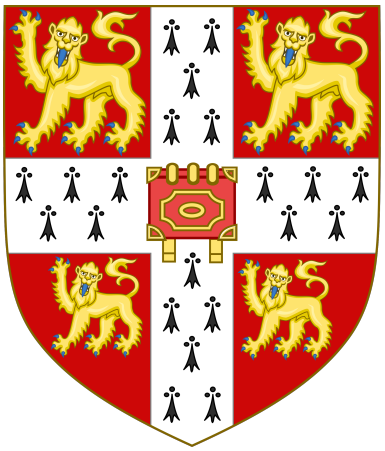
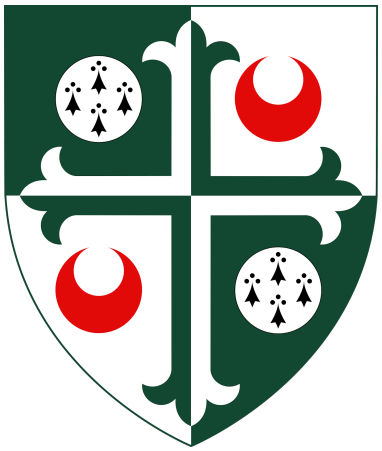
Data Analysis in Global 21-cm Experiments:
Chromatic Effects and Signal Reconstruction





Yu-Hsuan Shen

Girton College

&

Cavendish Laboratories
University of Cambridge

September 2025

This thesis is submitted for the degree of
Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

The research presented in this thesis is based on a series of first-author papers (Shen et al., 2022, 2024, 2025). Figure 1.4 is generated by the tool made available by Jiten Dhandha. The code used to generate the rotation measure synthesis technique cases in Chapter 4 is provided by Marta Spinelli.

Yu-Hsuan Shen
September 2025

Abstract

This thesis investigates chromatic effects that can distort the global 21-cm signal and introduces a flexible signal model designed to capture potential features arising from unexpected physics. The work is carried out within the framework of the REACH experiment and its data analysis pipeline.

Part I includes two chapters. Chapter 1 provides an introduction to 21-cm cosmology, and outlines both past and upcoming experimental efforts. Chapter 2 presents the numerical and statistical methods adopted throughout this thesis.

Part II and Part III contain the original research contributions, each chapter beginning with a contextual introduction.

Part II focuses on chromatic effects. Chapter 3, published as Shen et al. (2022), addresses the modelling of time-varying ionospheric effects and demonstrates that these effects do not average out over time, highlighting the importance of accounting for temporal variation in data analysis. Chapter 4, published as Shen et al. (2024), examines contamination from Galactic linear polarisation. Using three distinct models, the REACH pipeline is stress-tested and shown to recover injected global 21-cm signals in most cases, despite unaccounted-for polarisation.

In Part III, Chapter 5, the FlexKnot signal model is presented. This model is characterised by a twice-integrated cubic function built from freely positioned knots interpolated by piecewise linear splines. It is shown to effectively recover absorption features that may originate from unaccounted-for physics in the global 21-cm signal, performing robustly within the REACH data analysis framework. This work is published as Shen et al. (2025).

Finally, Part IV summarises the main conclusions and outlines potential directions for furthering these efforts.

Acknowledgements

If I had, in any way, made the course of this work appear to be less of an effort than it actually was, it would have only been the reflection of the people and circumstances that *under-stood* me along the way.

The first pots of my gratitude belong to my supervisors Eloy and Anastasia, who, besides giving academic directions and insights, have been more than *understanding* of my many hiccups. And, of course, Dominic, who has effectively been my supervisor from the start.

It has been an honour to be part of the REACH collaboration (and the funky office) for the exchange of both ideas and laughs. Special thanks to John for taking time to read and give meticulous comments on this very thesis.

The scholarship that underlies the entirety of my study was jointly offered by Cambridge Trust and Taiwan Ministry of Education. I also thank Girton College for the various forms of support they have kindly provided throughout my study.

This work would *most probably*[a] not have been possible without my parents, the progenitors of my existence and much of its ensuing maintenance that puts me in a position to do what I did. Credits also go to my sibling for never giving up on finding medical solutions to my funny predicament, even when I had.

I give my utmost appreciation to the many lending ears (and eyes), locally and across the ocean, for patiently enduring my endless thoughts at unheavenly hours on the ever-rolling woes and throes, and for offering advice and perspective with *understanding*—in particular, those belonging to Leann, Carrie, June, Carol, Natsumi, Yvonne, and Yifan. And now, a passing strange tribute to the mouth of Val, which hath, for many moons, boldly received the oft-times suspicious products of my culinary ventures (Gladkova, 2024).

Finally, I thank the Universe for its wonder. Thank Heavens!

[a] In that someone else might have done it anyway (but it has to be in the exact same way if one wants to be pedantic).

Contents

Declaration iii

Abstract v

Acknowledgements vii

Contents ix

Nomenclature xi

I Introduction 1

1 The Emergence of 21-cm Cosmology 3
- 1.1 The Elusive Period in Cosmic History . . . 3
- 1.2 21-cm Radio Cosmology . . . 4
- 1.3 Status of 21-cm Experiments . . . 10
- 1.4 Experimental Challenges . . . 11

2 Methods 17
- 2.1 Bayesian Statistics . . . 18
- 2.2 Nested Sampling . . . 19
- 2.3 PolyChord . . . 21
- 2.4 REACH Data Analysis Pipeline . . . 22

II Chromatic Effects and Foregrounds 27

3 Ionospheric Effects 29
- 3.1 Introduction . . . 29
- 3.2 Modelling Chromatic Ionospheric Effects . . . 33
- 3.3 Results . . . 36
- 3.4 Conclusions . . . 43

4 Foreground Polarisation 47
- 4.1 Introduction . . . 47
- 4.2 Simulating Observation with Linearly Polarised Foregrounds . . . 56
- 4.3 Results . . . 59
- 4.4 Conclusions . . . 66
- 4.A Difference between Two Orthogonally Pointing Radiometers . . . 69

III Global Signal Model **71**

5 FlexKnot Signal Model **73**
5.1 Introduction 73
5.2 Constructing the FlexKnot Function 75
5.3 Simulated Antenna Temperature Data 80
5.4 Results 84
5.5 Conclusions 93

IV Conclusions **95**

6 Conclusions **97**
6.1 Conclusions 97
6.2 Future 98

Nomenclature

Physical Constants

ϵ_0	permittivity of free space
c	speed of light
e	electron charge
h	dimensionless Hubble constant
$k_{\rm B}$	Boltzmann's constant
m_e	effective mass of the electron

Universal Notation

λ	wavelength
$\mathcal{L}$	likelihood
$\mathcal{Z}$	Bayesian evidence
ν	frequency
τ	optical depth
$N_{\rm e}$	electron density
T_α	brightness temperature of the Lyman-α radiation field
$T_{\rm CMB}$	brightness temperature of of the CMB
$T_{\rm f}$	brightness temperature of the foreground
$T_{\rm K}$	kinetic temperature of gas
$T_{\rm sky}$	brightness temperature of the sky model
$T_{\rm S}$	spin temperature of neutral hydrogen
T_{21}	brightness temperature of the global 21-cm signal
T_γ	brightness temperature of the radio background
z	redshift
$B(\nu, \theta, \phi)$	an antenna directivity pattern

Abbreviations

CMB	cosmic microwave background
EoR	epoch of reionisation
EUV	extreme ultraviolet
GSM	global sky model
HI	neutral hydrogen
IGM	intergalactic medium
LST	local sidereal times
RFI	radio frequency interference
RMSE	root-mean-square error
UV	ultraviolet

Project Abbreviations

AARTFAAC	Amsterdam-ASTRON Radio Transients Facility And Analysis Center
BIGHORNS	Broadband Instrument for Global HydrOgen ReioNisation Signal
CHIME	Canadian Hydrogen Intensity Mapping Experiment
DAPPER	Dark Ages Polarimeter PathfindER
DARE	Dark Ages Radio Explorer
DARELuSEE-Night	The Lunar Surface Electromagnetics Experiment at Night
DEX	Dark Ages Explorer
EDGES	Experiment to Detect the Global EoR Signature
GMRT	Giant Metrewave Radio Telescope
HERA	Hydrogen Epoch of Reionization Array
LARAF	Large-scale Array for Radio Astronomy on the Farside
LEDA	Large Aperture Experiment to Detect the Dark Ages
LOFAR	Low-Frequency Array
MIST	Mapper of the IGM Spin Temperature
MWA	Murchison Widefield Array
NenuFAR	New Extension in Nançay Upgrading LOFAR
OVRO-LWA	Owens Valley Radio Observatory Long Wavelength Array
PAPER	Precision Array for Probing the Epoch of Reionization

PRATUSH	Probing ReionizATion of the Universe using Signal from Hydrogen
PRIZM	Probing Radio Intensity at high-Z from Marion
REACH	Radio Experiment for the Analysis of Cosmic Hydrogen
RHINO	Remote HI eNvironment Observer
SARAS	Shaped Antenna measurement of the background Radio Spectrum
SCI-HI	Sonda Cosmológica de las Islas para la Detección de Hidrógeno Neutro
SKA	Square Kilometre Array radio telescope
SKA-Low	Square Kilometre Array Low frequency radio telescope

Part I

Introduction

Chapter 1

The Emergence of 21-cm Cosmology

The course of true line never did run smooth.

Pear W. Shakes

1.1 The Elusive Period in Cosmic History

Our understanding of the Universe—from its infancy, some 380,000 years after the big bang, to the present day—has matured significantly over the last several decades. This stems primarily from two major observational efforts: measurements of the cosmic microwave background (CMB) and galaxy surveys (Kellermann et al., 1969; Laureijs et al., 2011; Dey et al., 2019).

The CMB is the relic radiation emitted when the Universe cooled sufficiently for neutral atoms to form from the primordial plasma of nuclei and electrons, allowing the background radiation field to decouple from the baryonic matter, approximately 380,000 years after the Big Bang. First discovered by Penzias and Wilson in 1965 (Penzias and Wilson, 1965), subsequent microwave experiments, including ground-based (Keating et al., 2003; Kosowsky, 2003) and space-based ones (Smoot et al., 1992; Spergel et al., 2003; Planck Collaboration et al., 2020), have enabled increasingly detailed maps of the CMB. These observations reveal a nearly homogeneous and isotropic Universe whose baryonic matter is made up almost entirely of hydrogen and helium. It is a near-perfect black body with temperature perturbations at the level of one part in 10^5 (Bennett et al., 2003; Fixsen, 2009; Planck Collaboration et al., 2020). The power spectrum of these anisotrop-

ies is non-trivial and contains deep insights into the cosmological parameters of the Universe in its infancy.

In contrast, the present-day Universe is sprinkled with myriad structures; roughly 10% of baryonic matter has collapsed into stars and galaxies (Wang and Wei, 2023). Our understanding of the late-time Universe is mainly puzzled out from galaxy surveys, which have evolved from early visual catalogues to comprehensive, multi-dimensional mapping projects. The James Webb Space Telescope (Gardner et al., 2006) possesses the sensitivity to detect galaxies forming as early as a few hundred million years after the Big Bang (corresponding to redshifts $z \sim 10 - 20$). This technological leap has dramatically refined our ability to understand the growth and evolution of the galaxy population from the era just before reionisation.

Yet, between the first glow of the CMB and the dawn of the first lights that rekindled the Universe's darkness after recombination lies the elusive, a transformative period during which primordial density perturbations collapsed to form the first non-linear structures (Press and Schechter, 1974), nucleosynthesis seeded heavier elements (Suess and Urey, 1956), and the initially near-homogeneous Universe evolved into a state of increasing complexity under the influence of gravitational instability (Jeans, 1902). Gravity drove the initially small perturbations to grow into larger non-linear structures that eventually collapsed to become sheets, filaments and halos. What ensued were cycles of galaxy and star formations, and the transition from a neutral to an ionised Universe; a period we know pitifully little about outside theoretical modelling and prediction. This epoch is the missing thread in cosmic history, essential for weaving together our understanding of how the Universe evolved from apparent simplicity to the vast and intricate structures we observe today.

1.2 21-cm Radio Cosmology

Atomic hydrogen is the predominant form of baryonic matter in the interstellar medium between Recombination and the Epoch of Reionisation (EoR). Its hyperfine transition constitutes one of the very few transitions in the primordial gas that are sensitive to the thermal state of the intergalactic medium (IGM) during these early stages of the Universe (Furlanetto et al., 2006). Observing this hydrogen and its interaction with the background black-body radiation field offers a powerful tool for probing the cosmological and astrophysical phenomena during the dark ages and cosmic dawn. This forms the premise of 21-cm radio cosmology, which is to measure the redshifted 21-cm spectral line of primordial hydrogen in hope of gaining a better understanding of the infant Universe.

The 21-cm line of neutral hydrogen arises from the triplet-to-singlet hyperfine transition within the $1S$ ground state (Ewen and Purcell, 1951). This transition

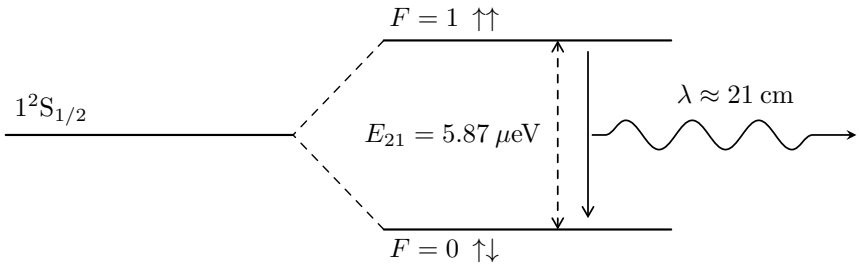


Figure 1.1: Diagram of the hyperfine transition (emission) of neutral atomic hydrogen that gives rise to the 21-cm signal. This transition is forbidden by the electric dipole selection and the difference between the energy levels, E_{21}, is low, resulting in a low spontaneous transition rate.

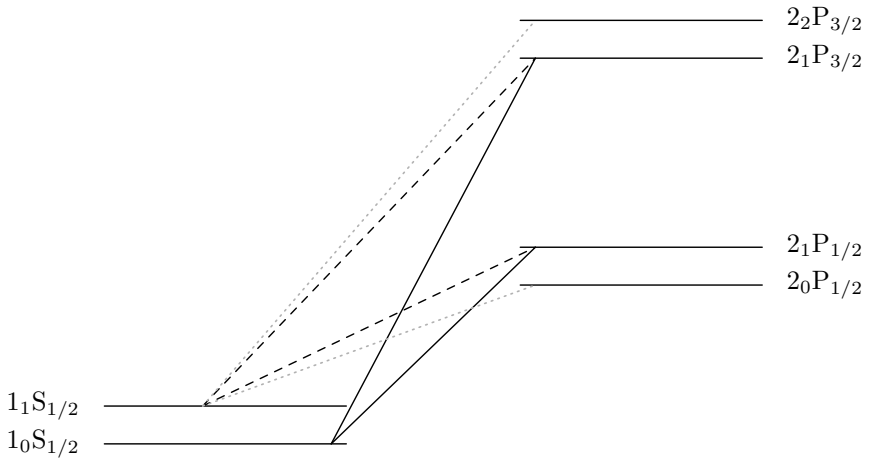


Figure 1.2: Diagram of the neutral hydrogen atom's hyperfine structure, showing the transitions relevant to the Wouthuysen-Field effect. Solid lines indicate allowed transitions by which electrons in the lower energy hyperfine level can be excited, while dashed lines show the corresponding de-excitations from the higher level. The gray dotted lines represent allowed transitions that do not contribute to the Wouthuysen-Field mechanism.

arises from the interaction between the magnetic moments of the proton and electron spins, resulting in the higher-energy triplet state with parallel spins and the lower-energy singlet state with antiparallel spins (see Figure 1.1). The hyperfine levels are separated by an energy difference of 5.87 μeV, corresponding to a photon with a frequency of approximately 1420 MHz, or a rest-frame wavelength of $\sim$ 21 cm. The transition between these two states is forbidden by the electric dipole selection, with a low transition probability of $2.85 \times 10^{-15} s^{-1}$.

The net rate of absorption or emission at this transition is determined by the number density of atoms in the lower n_0 and upper n_1 hyperfine energy state. The relative populations of the hyperfine states in terms of spin temperature, $T_{\rm S}$, can be defined using the Boltzmann distribution

$$\frac{n_1}{n_0} = \left(\frac{g_1}{g_0}\right) \exp\left(-\frac{E_{21}}{k_{\rm B} T_{\rm S}}\right), \tag{1.1}$$

where $E_{21} = 5.87 \mu$eV is the energy difference between the parallel and antiparallel energy states, g_i denotes the statistical weight, and $k_{\rm B}$ is the Boltzmann constant. The ratio $g_1/g_0 = 3$ indicates that the higher energy state is threefold degenerate.

The 21-cm signal becomes observable only when $T_{\rm S}$ at the 21-cm rest-frame wavelength deviates from the radio background, T_γ[a]. The signal's magnitude depends on the hydrogen neutral fraction, $x_{\rm HI}$, the number density of hydrogen, $n_{\rm H}$, and the redshifting of photons as the Universe expands. The optical depth of a cloud of hydrogen, τ_ν, can be derived by integrating the absorption coefficient, α_ν, along the path through the cloud:

$$\begin{aligned} \tau_\nu = \int \alpha_\nu \, {\rm d}s &= \int -\frac{E_{21}}{4\pi} \phi(\nu)(n_0 B_{10} - n_1 B_{01}) \, {\rm d}s \\ &= \int \left[1 - \exp\left(-\frac{E_{21}}{k_{\rm B} T_{\rm S}}\right)\right] \sigma_0 \phi(\nu) n_0 \, {\rm d}s, \end{aligned} \tag{1.2}$$

where $n_0 = n_{\rm H}/4$, $\phi(\nu)$ is the line profile function, and $\sigma_0 \phi(\nu)$ is the 21-cm cross section. Here $\sigma_0 \equiv 3c^2 A_{10}/8\pi\nu^2$, with $A_{10} = 2.85 \times 10^{-15} {\rm s}^{-1}$, the spontaneous emission rate of the spin-flip transition. The resulting integral of τ_ν is:

$$\begin{aligned} \tau_\nu &= \frac{3}{32\pi} \frac{E_{21}}{k_{\rm B} T_{\rm S}} \frac{A_{10}}{\nu^3} \frac{1}{1+z} \frac{x_{\rm HI} n_{\rm H}}{{\rm d}v_\parallel / {\rm d}r_\parallel} \\ &\approx 0.0092 (1+\delta)(1+z)^{\frac{3}{2}} \frac{x_{\rm HI}}{T_{\rm S}} \frac{H(z)/(1+z)}{{\rm d}v_\parallel / {\rm d}r_\parallel}. \end{aligned} \tag{1.3}$$

For a given redshift, z, the differential 21-cm brightness temperature, δT, for a current observer is:

$$\delta T(z) = (1 - e^{-\tau_\nu}) \frac{T_{\rm S} - T_\gamma(z)}{1+z}. \tag{1.4}$$

The exponential can be Taylor expanded:

$$\begin{aligned} \delta T(z) &\approx \frac{T_{\rm S} - T_\gamma(z)}{1+z} \tau_\nu \\ &\approx 9 x_{\rm HI} (1+\delta)(1+z)^{\frac{1}{2}} \left[1 - \frac{T_\gamma(z)}{T_{\rm S}}\right] \frac{H(z)/(1+z)}{{\rm d}v_\parallel / {\rm d}r_\parallel} \, {\rm mK}. \end{aligned} \tag{1.5}$$

[a] T_γ is usually dominated by the CMB ($T_{\rm CMB}$), so $T_{\rm CMB}$ is adopted instead in some literature.

Equation (1.5) shows that δT reaches saturation when $T_S \gg T_\gamma$, and can become very large when $T_S \ll T_\gamma$, indicating the dependence of the brightness temperature on T_S.

The spin temperature of hydrogen is determined by the relative populations of atoms in the excited and ground states, and is influenced by three processes:

(i) **CMB absorption and emission**[b]**:** Hydrogen atoms absorb and emit 21-cm photons, leading to spin-flip transitions. This process drives T_S towards thermal equilibrium with the CMB, such that $T_S \to T_{CMB}$. At equilibrium, the net spin-flip rate becomes zero, and neutral hydrogen fully couples with the CMB photons. Without additional processes to break this coupling, T_S would rapidly reach equilibrium, resulting in null detectable 21-cm signal, $\delta T = 0$ (Furlanetto et al., 2006).

(ii) **Collisional coupling:** Collisions with other hydrogen atoms or free electrons tend to couple T_S to the kinetic temperature of the gas, T_K. This process is efficient at high redshifts, $z \sim 70$, where the gas density is sufficient for frequent collisions. However, by $z \sim 30$, the expansion of the universe reduces gas density, rendering collisions ineffective in altering T_S. At this stage, CMB coupling dominates except in very dense minihalos, where collisions remain significant (Loeb and Zaldarriaga, 2004; Hirata, 2006).

(iii) **Wouthuysen-Field mechanism:** The scattering of Lyman-α photons, through the Wouthuysen-Field mechanism (Field, 1958; Wouthuysen, 1952), plays a critical role in decoupling T_S from the CMB. Although direct transitions between the $n = 1$ hyperfine levels of neutral hydrogen are forbidden, Lyman-α photons enable transitions via the allowed excitation to $n = 2$ states (see Figure 1.2). This process redistributes electrons asymmetrically between the hyperfine levels, driving T_S towards T_α, the brightness temperature of the Lyman-α radiation field (Pritchard and Furlanetto, 2006). In the optically thick intergalactic medium at high redshifts, repeated scattering ensures that $T_\alpha \to T_K$, thereby coupling T_S to T_K (Furlanetto, 2016).

The rate of these processes are fast compared to the de-excitation of the line, allowing the spin temperature to be approximated given the equilibrium balance of these effects. The efficiency of these processes is represented by their respective coupling coefficients x_γ, x_α, and x_c, and the latter two are determined relative to the rate due to the radio background: x_α can be calculated using molecular potentials and multichannel quantum collision theory (Zygelman, 2005) and x_c by the intensity of Lyman line radiation fields (Hirata, 2006). The spin temperature can then be written as a linear combination of these effects, weighted by their

[b] Some studies (Venumadhav et al., 2018) suggest that the Lyman-α photons also mediate energy transfer between the CMB photons and the thermal motions of the hydrogen atoms in the IGM. This is referred to as CMB heating.

respective coefficients:

$$T_{\rm S}^{-1} = \frac{x_\gamma T_\gamma^{-1} + x_\alpha T_{\rm c}^{-1} + x_{\rm c} T_{\rm K}^{-1}}{x_\gamma + x_\alpha + x_{\rm c}}, \tag{1.6}$$

where $T_{\rm c}$ is the colour temperature of Lyman-α photons (Barkana, 2016).

Although the transition has a rest-frame wavelength of 21 cm when it is produced, the expansion of the Universe significantly redshifts it by the time it reaches us, resulting in the observed wavelength of $\lambda = 21(1+z)$ cm. This redshifting enables us to map the 21-cm signal across cosmic time using frequency.

1.2.1 Evolution of the redshifted 21-cm Signal

Neutral atomic hydrogen (Figure 1.3) first forms during recombination at $z \sim 1100$. At these high redshifts, gas is sufficiently dense that Compton scattering transfers energy from scattering ultraviolet CMB photons to residual free electrons, driving $T_{\rm K}$ to T_γ. Simultaneously, collisions effectively couple $T_{\rm K}$ with $T_{\rm S}$, causing $T_{\rm S}$ to remain just slightly lower than T_γ due to adiabatic cooling. This produces a subtle absorption feature against the radio background in the sky-averaged, or global, 21-cm signal.

At $z \sim 300$, the Compton scattering starts to become inefficient and by $z \sim 150$ negligible, as gas density decreases due to the ongoing expansion of the Universe. $T_{\rm K}$ decreases adiabatically as $(1+z)^2$, and $T_{\rm S}$ decreases similarly while collisional coupling remains efficient. On the other hand, T_γ decreases more slowly, following a $1+z$ dependence, a consequence of its wavelengths redshifting with cosmic expansion. The growing disparity between the two temperatures results in a progressively stronger absorption feature in the sky-averaged 21-cm signal over time during this period.

Due to the expansion of the Universe, the gas cools and becomes diluted, rendering collisional coupling increasingly inefficient. At around $z \sim 80$, the coupling of $T_{\rm S}$ to T_γ begins to gradually dominate over the rapidly decreasing coupling to $T_{\rm K}$ via collisions. Eventually $T_{\rm S}$ tends back to approach the radio background temperature, leading to the near extinction of the 21-cm signal at $z \sim 30$ (Furlanetto et al., 2006). Collisional coupling introduces a peak in the 21-cm power spectrum during these high redshifts, specifically from the scales corresponding to the mean free path for collisions, where the coupling between gas kinetic temperature and spin temperature is most effective.

The first stars begin to form at $z \sim 65 - 25$ (Fialkov et al., 2012), providing the conditions for the Wouthuysen-Field effect (Wouthuysen, 1952; Field, 1958) to significantly influence the coupling of spin temperature to the gas kinetic temperature via Lyman-band photons emitted by these luminous objects. The first stars illuminate the intergalactic medium with their Lyman-band emission from $z \sim 65$ onwards, with its intensity depending on the star formation rate in early galaxies.

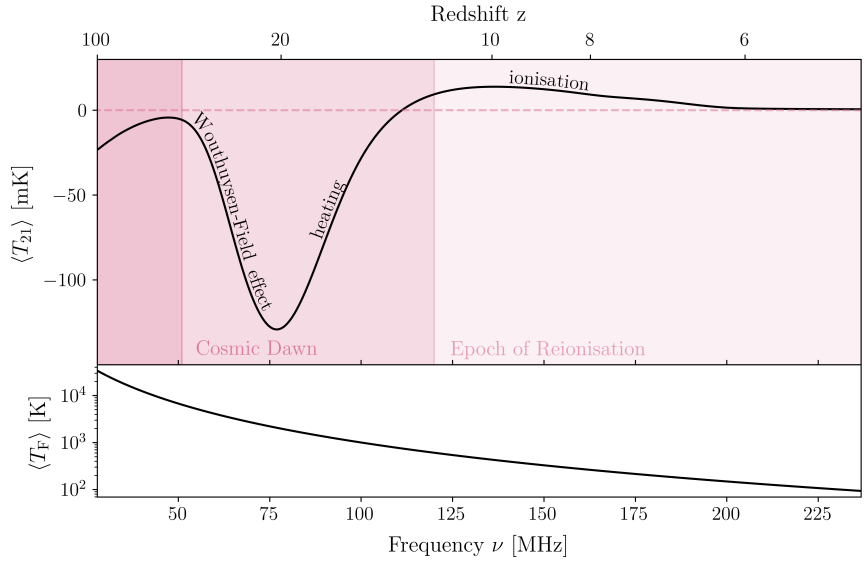


Figure 1.3: The upper panel is a simulated global 21-cm signal $\langle T_{21} \rangle$ generated using GLOBALEMU (Bevins et al., 2021) with the following parameters: star formation efficiency $f_* = 0.02$, minimal virial circular velocity $V_c = 16.5$ km/s, X-ray efficiency $f_x = 1$, CMB optical depth $\tau = 0.06$, slope of the spectral energy density $\alpha = 1$, low energy cut-off of the X-ray spectral energy density $\nu_{min} = 0.2$ keV, mean free path of the ionising photons $R_{mfp} = 30$ Mpc. The bottom panel is a simulated sky-averaged astrophysical foreground $\langle T_F \rangle$ in log scale, which is considered smooth compared to the global 21-cm signal. The 21-cm signal first emerges during the cosmic Dark Ages, undergoes significant evolution during Cosmic Dawn, and then effectively disappears from the intergalactic medium as ionised bubbles coalesce and complete the Epoch of Reionisation (EoR).

This causes the spin temperature to fall back towards the cold gas temperature, resulting in a pronounced absorption feature against the radio background in the 21-cm signal.

During the emergence of numerous luminous sources, additional heating processes, including Lyman-α and X-ray heating, become effective in raising the temperature of the hydrogen gas. Alongside the contribution from shock heating, these mechanisms counteract the adiabatic cooling, which would otherwise be faster than the cooling of the radio background. The coupled T_S and T_K gradually increase towards the radio background temperature T_γ. In some models, they may even exceed T_γ, producing an emission feature by $z \sim 10$ (Muñoz, 2023).

The luminous sources emit a substantial number of high-energy ultraviolet photons capable of ionising atomic hydrogen in the intergalactic medium, thereby creating bubbles of fully ionised hydrogen around these sources. Initially, these ionised bubbles form individually and remain relatively small. Within these bubbles, the neutral hydrogen fraction is zero, resulting in the absence of a 21-cm signal. This marks the beginning of the EoR. Over time, as more luminous sources emerge

and their radiation intensifies, the ionised regions expand and begin to merge. Eventually, the ionised regions encompass the entire intergalactic medium, culminating in the near-complete[c] ionisation of hydrogen and the near-extinction[d] of the cosmological 21-cm signal by $z \sim 5.5$ (Shaver et al., 1999; Furlanetto, 2016). At this point, reionisation is complete.

The overall structure of the expected 21-cm signal is well understood, but the exact evolution is dependent on the timing, amplitude, and rate of crucial astrophysical events in cosmic history. Its recovery could provide insights into the properties and formation of the first stars and galaxies. Additionally, the 21-cm signal offers a way to verify potential exotic physics, such as cold dark matter, which could modify its expected evolution and, in turn, reveal the nature of dark matter. This vast potential to illuminate the astrophysical and cosmological processes in the early Universe motivated the current efforts to measure it.

1.3 Status of 21-cm Experiments

Currently planned and ongoing efforts to detect the 21-cm signal are either global experiments or full power spectrum measurements. Global experiments employ one single broadband antenna with a wide beam (Shaver et al., 1999), aiming to detect the sky-averaged, or global, 21-cm signal, as a function of redshift, z. Historical examples include SCI-HI (Voytek et al., 2014), BIGHORNS (Sokolowski, 2015), EDGES 2 (Bowman et al., 2018), SARAS 2 (Singh et al., 2018), and LEDA (Price et al., 2018), while ongoing experimental efforts include SARAS 3 and 4 (Singh et al., 2022), REACH (de Lera Acedo et al., 2022), EDGES 3 (Cappallo et al., 2023), PRIZM (Philip et al., 2019), MIST (Monsalve et al., 2024), and RHINO (Bull et al., 2024). In addition to these Earth-based radiometers, proposals for space-based global 21-cm signal experiments include DARE (Burns et al., 2012), FARSIDE (Hallinan et al., 2021), DSL (Shi et al., 2022), PRATUSH (Sathyanarayana Rao et al., 2023), LuSEE-Night (Bale et al., 2023), FarView (Polidan et al., 2024), LARAF (Chen et al., 2024), CosmoCube (Artuc and de Lera Acedo, 2025), and ALO/DEX (Brinkerink et al., 2025).

The work presented in this thesis has been conducted as part of one of these ongoing efforts: the Radio Experiment for the Analysis of Cosmic Hydrogen (REACH). The models and analysis techniques developed herein contribute to the foundational framework for this experiment.

On the other hand, experiments that attempt to measure the full 21-cm power spectrum include OVRO-LWA (Ellingson et al., 2009), MWA (Lonsdale et al.,

[c] There remain pockets of neutral gas inside galaxies protected from the ionising radiation as they are located in very dense regions—the process is known as self-shielding (Chardin et al., 2018).

[d] The self-shielded regions are present all the way to $z \sim 0$ and produce 21-cm signal that can be observed via intensity mapping.

2009), PAPER (Parsons et al., 2010), GMRT (Chang and GMRT-EoR Team, 2013), LOFAR (van Haarlem et al., 2013), HERA (DeBoer et al., 2017), LEDA (Price et al., 2018), AARTFAAC (Gehlot et al., 2020), and NenuFAR (Mertens et al., 2021). Power spectrum measurement is achieved by correlating signals between antenna pairs, which measures the Fourier modes of the sky. By combining the spatial scales from baselines with the depth scales from frequency bandwidth, the variance of the 21-cm brightness temperature fluctuations is calculated across the observed volume. No detection has been reported by this approach, but upper limits on the 21-cm power spectrum have been suggested (Mertens et al., 2025; Nunhokee et al., 2025; Munshi et al., 2025). Figure 1.4 shows the compilation of several published upper limits. By far, the most sensitive upper limits on the magnitude of the 21-cm power spectrum are given by HERA (HERA Collaboration et al., 2023) at $z \sim 7.9$ and $z \sim 10.4$. These results paved the way for the next generation interferometers such as the SKA-Low, which will be capable of creating full three-dimensional tomographic maps of the 21-cm signal.

In 2018, Bowman et al. (2018) reported a sky-averaged 21-cm signal detection in the shape of a flattened Gaussian centred at 78 MHz with an amplitude of 500 mK. The amplitude of the signal exceeds expectations from standard astrophysical models. The signal's unusual flat bottom and steep sides are also challenging to replicate in theoretical models. Exotic astrophysical explanations have been proposed, including the thermal coupling of dark matter to baryonic matter (Barkana, 2018; Muñoz et al., 2018) and an excess radio background from high-redshift galaxies or other exotic sources (Fialkov and Barkana, 2019; Reis et al., 2020). Re-examinations on their methods, however, have raised concerns regarding the non-physical parameters, the non-uniqueness of their solution, and potentially unaccounted-for systematic structures either in the data or the analysis (Hills et al., 2018; Singh and Subrahmanyan, 2019). Furthermore, Hills et al. (2018), Sims and Pober (2020), and Bevins et al. (2020) demonstrated that a damped sinusoidal systematic is strongly favoured in the EDGES data. The SARAS collaboration also reported non-detection of the best-fit EDGES profile with 95.3% confidence using their SARAS 3 instrument (Singh et al., 2022).

1.4 Experimental Challenges

As established, the 21-cm signal can provide crucial insights into the processes that shaped the infant Universe. However, the practical task of measuring this signal is exceptionally difficult. Its brightness temperature is only a few millikelvin, making it susceptible to numerous larger contaminants. The following discussion will focus specifically on the challenges confronting global experiments, which are the primary concern of this thesis.

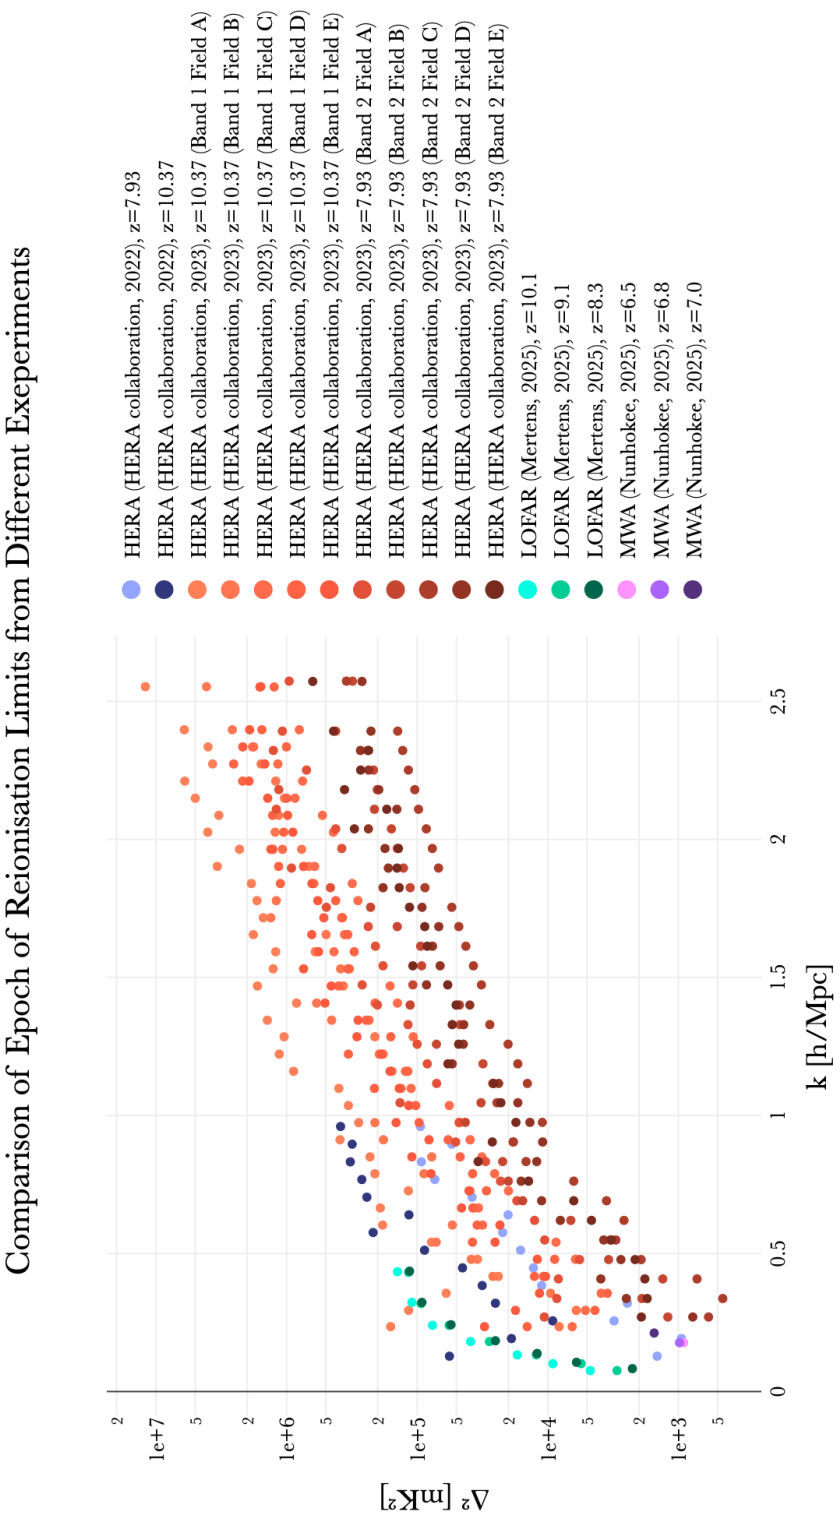


Figure 1.4: A compilation of currently published EoR limits from various full 21-cm power spectrum experiments. This figure was generated using the publicly available `EoRImaging` tool, developed by Steven Murray and redesgined by Jiten Dhandha.

1.4.1 Astrophysical Foregrounds and Beam Chromaticity

When radiometers measure incoming radiation, they inevitably pick up unwanted emissions from sources including our Galaxy and extra-galactic ones; together they introduce a foreground that obscures the desired 21-cm signal. This foreground consists predominantly of diffuse synchrotron radiation and free-free emission from the Milky Way, along with some point source emission from other galaxies. These emissions are expected to follow an approximate power law of $\nu^{-2.5}$, where ν represents the frequency. This is considered smooth compared to the expected shape of the 21-cm signal (see Figure 1.3). Within the $50 - 200$ MHz frequency band, which global 21-cm experiments operate within, the foreground emissions, often thousands of Kelvin, are about 5 orders of magnitude larger than the 21-cm signal, which is expected to be only tens to hundreds of millikelvin. Therefore, separating the 21-cm signal from such bright foreground emissions is a non-trivial task.

For full power spectrum measurements, a technique called 'foreground avoidance' (Chapman et al., 2016) can be adopted. In this approach, the smoothness of the foregrounds confines their contamination to a specific region of Fourier space, the so-called 'foreground wedge', leaving other regions, in principle, clean for a measurement. This strategy, however, cannot be applied to global experiments because they measure the total sky-averaged signal without spatial resolution, integrating over all lines of sight. Consequently, global experiments often rely on a different strategy known as 'foreground removal', which utilises the spectral smoothness of the foregrounds by modelling them with flexible functions, such as power laws, log-polynomials, or maximally smooth functions, and subtracting them from the data.

This is further complicated when the foregrounds are convolved with the antenna beam. A wide beam is necessary to measure the sky-averaged 21-cm signal because its evolution spans a broad frequency range. However, a wide bandwidth allows for a larger number of current modes in the antenna, increasing the potential for variations in the currents induced by incident waves. This introduces very non-smooth features that could be challenging to correct for in the originally smooth foregrounds when they are convolved. This is often referred to as 'beam chromaticity', and it poses a major challenge for 21-cm experiments (Tauscher et al., 2020).

The real sky is non-homogeneous and rich in structure, with various sources of radiation. When the frequency-dependent antenna beam pattern is coupled with this complex spatial variation from the sky, it introduces distortions to the observed foregrounds in the frequency domain. One approach to correct for this involves simulating the gain pattern and convolve it with a sky model scaled at a

reference frequency $T_{\rm sky}(\nu_{\rm ref})$. The correction[e] is usually expressed as a variation on

$$B(\nu,t) = \frac{\int_\Omega T_{\rm sky}(\nu,\Omega,t)D(\nu,\Omega)d\Omega}{\int_\Omega T_{\rm sky}(\nu_{\rm ref},\Omega,t)D(\nu_{\rm ref},\Omega)d\Omega}, \tag{1.7}$$

where

$$T_{\rm sky}(\nu_{\rm ref},\Omega,t) = [T_{\rm map}(\Omega,t) - T_\gamma]\left(\frac{\nu_{\rm ref}}{\nu_{\rm map}}\right)^{-\beta} + T_\gamma, \tag{1.8}$$

$T_{\rm map}$ represents a base all-sky map at frequency $\nu_{\rm map}$ rotated to time t, $\nu_{\rm ref}$ is the reference frequency within the observed band, β is a model of the foreground spectral index, and D is the directivity pattern of the antenna beam. The correction is applied by dividing the observed data at each local sidereal time (LST) by the corresponding $B(\nu,t)$.

In addition, Anstey et al. (2021) proposed a sophisticated forward-modelling approach for the antenna beam and the foreground, where a base all-sky map is divided into N regions of varying spectral indices and then convolved with the beam model:

$$T_{\rm f} = \frac{1}{4\pi}\int_\Omega D(\nu,\Omega)\int_{t_0}^{t_1}\sum_{i=0}^{N} M_i(\Omega)(T_{\rm map} - T_\gamma)\left(\frac{\nu}{\nu_{\rm map}}\right)^{-\beta_i} dtd\Omega + T_\gamma, \tag{1.9}$$

where M_i is the mask that divides the sky map into N regions, each with a spectral index β_i. This is the method adopted by the REACH data analysis pipeline.

1.4.2 The Ionosphere

The ionosphere poses another challenge for 21-cm experiments. Extending from 50 km to beyond 1000 km above the Earth's surface, the ionosphere is a layer of magnetised plasma that distorts and attenuates incoming radiation. The effects scale as the square of its wavelength, making them particularly significant within the frequency range relevant to 21-cm observations.

The ionosphere is typically divided into three primary layers, D, E, and F, each characterised by distinct electron densities and ionic compositions. For 21-cm experiments, the most critical ionospheric effects arise from the D and F layers. The D layer introduces pronounced chromatic absorption due to its high neutral density, which results in a substantial electron collision frequency. The absorption attenuates the intensity of the observed signal. In contrast, the F layer primarily contributes to chromatic refraction owing to its high electron density. It results in the decrease in the zenith angle of the incoming ray so that they appear to the antenna at higher elevations than they are originally if the ionosphere was not present. Both of these effects can be analysed by the complex refractive index described by the Appleton-Hartree equation (Appleton, 1932).

[e] While used here as a foreground correction, this formalism can also be used more generally to quantify the chromaticity of an antenna's gain pattern.

There are few existing measurements of the ionosphere at night, particularly for the D layer, owing to its low electron density (Schunk and Nagy, 2000). Without accurate knowledge of the parameters that characterise these layers, the ionospheric effects cannot be precisely determined. Moreover, modelling these effects is challenging because ionospheric conditions vary significantly with location, time, season, the solar cycle, and magnetic activity. The chromatic ionospheric effects will be discussed in greater detail in Chapter 3.

1.4.3 Galactic Foreground Polarisation

Polarised foreground emission undergoes Faraday rotation as it passes through the magnetic fields of the interstellar medium, imprinting a chromatic structure that complicates the extraction of the 21-cm signal (Rybicki and Lightman, 1986). In experiments employing polarisation sensitive antennas, such as dipoles, the instruments inherently measure linearly polarised signals. Inadequate calibration can lead to leakage of the complex polarised signal into the total intensity.

Observations of linear Galactic polarisation at the relevant low frequencies remain limited, leaving its precise structure poorly understood. It is particularly troublesome because the effect is frequency-dependent: the change in the polarisation angle is proportional to the square of the radiation's wavelength. Consequently, leakage into the total intensity varies across frequency, making it challenging to disentangle from the unpolarised signal without substantial signal loss. Most importantly, the spectral structure of the Galactic polarised emission leakage can mimic the 21-cm signal. The impact of Galactic foreground polarisation contamination will be examined in detail in Chapter 4.

1.4.4 Others

There are myriad challenges to measuring the 21-cm signal on Earth. An example is radio frequency interference (RFI) (Ford and Buch, 2014; Baan, 2019) from various sources such as communication devices, radar systems, and satellites (Witze, 2022). The 21-cm signal, unfortunately, lies in unprotected frequency bands. RFI can diminish the sensitivity of a radio telescope and introduce artefacts into the observed data. Transient RFI, which is not constant over time, is particularly difficult to detect and avoid in large data sets with many time bins, such as those from 21-cm experiments. Anstey and Leeney (2024) present a method based on Leeney et al. (2023) to mitigate the interference using a Bayesian inference method, enabling transient anomalies to be flagged and accounted for in a Bayesian manner when fitting models to time-series data through likelihood re-weighting.

Other challenges include the reflective properties of the soil near the horizon (Pattison et al., 2024) and artefacts introduced by ground planes due to their finite edges and structural discontinuities (Bradley et al., 2019). The impact of mutual

coupling between instruments is also a significant challenge (O'Hara et al., 2025). Unknown systematics, arising from limited knowledge or incorrect handling, are likely to have significant impacts (Cumner et al., 2024). Separating the various potential causes of such systematics is difficult due to differing experimental setups, environmental conditions, and the time-varying nature of the data.

Chapter 2

Methods

Though this be madness, yet there is method in't.

W. Shakespeare

Finding ways to accurately describe the probabilities of different instances is quintessential in modern cosmology. In the Bayesian language, a probability quantifies the degree of confidence in a proposition or a model. The frequentist approach, on the other hand, defines probability as the relative frequency of an event occurring in a large number of repeated independent trials under identical conditions.

When the goal is to assess confidence in a model, the frequentist interpretation becomes limiting; it can only provide the probability of observing certain outcomes, not the statistical confidence in a hypothesis itself. This often leads to misinterpretation, where frequentist probabilities are mistakenly taken as measures of belief in a proposition. For these reasons, the Bayesian framework is more naturally suited for cosmological inference, where statistical confidence in competing hypotheses can be clearly and intuitively defined.

Bayesian inference is the primary statistical tool used in the REACH experiment. The REACH data analysis pipeline adopts a Bayesian nested sampling algorithm for model fitting. In this chapter, one can find an introduction to the Bayesian statistics and the REACH data analysis pipeline, which is mentioned throughout this thesis.

2.1 Bayesian Statistics

Bayes' theorem is encapsulated by the equation

$$\mathrm{P}(A|B) = \frac{\mathrm{P}(B|A)\mathrm{P}(A)}{\mathrm{P}(B)}, \tag{2.1}$$

where $\mathrm{P}(A|B)$ represents the conditional probability of event A given event B, provided that $\mathrm{P}(B) \neq 0$, the theorem demonstrates how conditional probabilities can be inverted, in this case by de-conditioning through multiplication by $\mathrm{P}(A)$ and reconditioning through normalisation by $\mathrm{P}(B)$. Now, if we substitute A with data, $\mathcal{D}$ and B with hypothesis, $\mathcal{H}$, it essentially means that we are now capable of, thanks to Bayes' theorem, deriving the probability of a hypothesis given the data $\mathrm{P}(\mathcal{H}|\mathcal{D})$.

2.1.1 Bayesian Inference

Bayesian inference is a statistical technique grounded in Bayes' theorem. It relates the probability of a model, given a set of data, $\mathcal{D}$, to the probability distribution of the data given the model, $\mathcal{M}$, conditioned on a set of parameters, $\theta_{\mathcal{M}}$, referred to as the *likelihood*, $\mathcal{L}$, which is typically easier to compute:

$$\mathcal{L} \equiv \mathrm{P}(\mathcal{D}|\theta_{\mathcal{M}}, \mathcal{M}). \tag{2.2}$$

An initial assumption about the parameters must be specified to perform Bayesian inference:

$$\pi \equiv \mathrm{P}(\theta_{\mathcal{M}}|\mathcal{M}). \tag{2.3}$$

This is known as the *prior distribution* of the parameter. It is common to choose a uniform or Gaussian prior on the parameter values, or log-uniform when the order of magnitude of the parameter is not well-known. The range of the initial prior distribution for a parameter should cover the expected scale and spread of its values, based on existing theoretical models and previous experimental results.

Model comparison

For a scientific problem, it is often the case that more than one model can be formulated to explain a particular dataset, and the problem now lies in how to determine which of these models should be favoured based on the data. One merit of Bayesian inference is its ability to perform model comparisons while simultaneously estimating parameters. The Bayesian 'evidence' $\mathcal{Z}$ quantifies how well a model explains the observed data. Often referred to as the 'marginal likelihood', it is computed by integrating the likelihood over the entire parameter space, weighted by the prior (Sivia and Skilling, 2006):

$$\mathcal{Z} \equiv \mathrm{P}(\mathcal{D}|\mathcal{M}) = \int \mathrm{P}(\mathcal{D}|\theta_{\mathcal{M}}, \mathcal{M})\mathrm{P}(\theta_{\mathcal{M}}|\mathcal{M})\,\mathrm{d}\theta_{\mathcal{M}} = \int \mathcal{L}\pi\,\mathrm{d}\theta_{\mathcal{M}}. \tag{2.4}$$

The evidence gives the probability of observing the data $\mathcal{D}$ conditioned on the model $\mathcal{M}$. What we actually desire is the probability of a model $\mathcal{M}$ given the data $\mathcal{D}$, which, fortunately, can be easily obtained using Bayes' theorem (Equation 2.1):

$$\mathrm{P}(\mathcal{M}|\mathcal{D}) = \frac{\mathrm{P}(\mathcal{D}|\mathcal{M})\mathrm{P}(\mathcal{M})}{\mathrm{P}(\mathcal{D})} = \mathcal{Z}\frac{\mathrm{P}(\mathcal{M})}{\mathrm{P}(\mathcal{D})}, \tag{2.5}$$

where $\mathrm{P}(\mathcal{D})$ serves as a normalisation factor independent of the model. The relative probabilities of competing models are given by the ratio of their evidences, weighted by their prior probabilities $\mathrm{P}(\mathcal{M})$.

A useful consequence of how the evidence is calculated (Equation 2.4) by integrating the prior-weighted likelihood over the entire parameter space is that models with more parameters tend to yield lower evidence values compared to models of similar performance but with fewer parameters. This down-weights models with more parameters or wider priors unless these features significantly improve the likelihood over their counterparts with fewer parameters or narrower priors. This feature inherently penalises overly complex models, favouring those that offer a better balance between explanatory power and simplicity.

Parameter estimation

Apart from selecting the best model, knowing what the data reveal about its parameters is of equal importance. Bayes' theorem, again, enables us to find the posterior distribution $\mathcal{P}$, which represents the updated knowledge of parameters $\theta_{\mathcal{M}}$ after incorporating the data $\mathcal{D}$, by inverting the conditioning in terms of likelihood, prior, and evidence:

$$\mathcal{P} \equiv \mathrm{P}(\theta_{\mathcal{M}}|\mathcal{D},\mathcal{M}) = \frac{\mathrm{P}(\mathcal{D}|\theta_{\mathcal{M}},\mathcal{M})\mathrm{P}(\theta_{\mathcal{M}}|\mathcal{M})}{\mathrm{P}(\mathcal{D}|\mathcal{M})}, \quad \text{or} \quad \mathcal{P} = \frac{\mathcal{L}\times\pi}{\mathcal{Z}}. \tag{2.6}$$

This equation illustrates how our initial knowledge of the parameters π updates to posterior $\mathcal{P}$ based on the data $\mathcal{D}$.

2.2 Nested Sampling

Nested sampling is a method devised by Skilling (2006) to efficiently calculate the evidence (Equation 2.4). The primary function of nested sampling is to calculate the Bayesian evidence, thereby enabling model comparison, while parameter estimation serves as a by-product. The integral used to calculate the evidence is typically over a high-dimensional parameter space, but only a small fraction of this space will contribute to the evidence. This small space will be the region surrounding the peak of the posterior, and it is challenging to locate in a high dimensional space.

2.2.1 Prior Volume Transformation

Nested sampling works around the difficulty of calculating high dimensional evidence integral by making the prior volume X represent a fraction of the total prior volume within an iso-likelihood contour. A likelihood contour is defined as

$$C(\mathcal{L}_*) = \{\theta : \mathcal{L}(\theta) = \mathcal{L}_*\}. \tag{2.7}$$

Prior volume X is a region containing some fraction of the prior enclosed by this contour of likelihood $\mathcal{L}_*$:

$$X(\mathcal{L}) = \int_{\mathcal{L}(\theta)>\mathcal{L}_*} \pi(\theta)\,\mathrm{d}\theta. \tag{2.8}$$

In other words, for each likelihood contour, there is a corresponding likelihood and a prior volume. This reformulation effectively reduces the calculation to a one-dimensional integral.

2.2.2 Multi-point Nested Sampling

The nested sampling algorithm begins by uniformly sampling live points n_{live} from the prior distribution of the parameter $\pi(\theta)$ and evaluating their likelihoods. In each iteration, the point with the lowest likelihood is discarded and replaced by a new point that has a higher likelihood. The new point is drawn uniformly[a] from the prior distribution. By this process, the prior volume of a parameter contained by the live points would contract by a fixed fraction, or exponentially, towards the peak of the posterior:

$$\langle X_i \rangle \approx e^{-\frac{i}{n_{\text{live}}}}, \tag{2.9}$$

with the initial value normalised to $X_0 = 1$.

The prior can be expressed as the gradient of the prior volume with respect to the parameters θ:

$$\pi(\theta) = \frac{\mathrm{d}X}{\mathrm{d}\theta}. \tag{2.10}$$

The evidence (Equation 2.4) can then be approximated using the weighted likelihoods of the discarded (or dead) points and the final remaining (live) points, $\mathcal{L}_i^*$:

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)\,\mathrm{d}X \approx \sum_i^{n_{\text{iter}}} w_i \mathcal{L}_i^*, \tag{2.11}$$

where n_{iter} is the number of iterations. The weights w_i are given by $w_i = X_{i+1} - X_i$, representing the difference in prior volume between successive iteration. It is worth noting that, for many parameters in cosmological problems, a significant portion of the posterior likelihood will occupy only a small fraction of the prior space. Taking this into account, it is often preferable to take samples uniformly in $\log X$

[a] *Uniformly* means every valid point within the prior region has an equal chance of being the next replacement.

space rather than X space to effectively sample the prior space. Eventually, the posterior distribution is represented by all the likelihood samples weighted by priors, clustering around the parameter values with the highest likelihood, where most of the posterior mass is concentrated.

2.2.3 Terminating the Iteration

As nested sampling proceeds, the likelihood value $\mathcal{L}_i$ monotonically increases, while the weight w_i monotonically decreases. In the beginning, $\mathcal{L}_i$ increases faster than w_i decreases. At some point, $\mathcal{L}_i$ will reach a point where it flattens sufficiently that the decreasing w_i then dominates. Consequently, the function $X\mathcal{L}(X)$ reaches a peak. The algorithm is set to terminate when the remaining posterior mass left within the live points becomes some fraction of the current evidence:

$$X_i \times \langle \mathcal{L} \rangle_{\text{live}} \ll \mathcal{Z}. \tag{2.12}$$

In high-dimension spaces, the absolute peak of the likelihood only occupies a negligible fraction of the prior mass, making it unnecessary to identify the maximum likelihood for the purposes of inference.

2.3 PolyChord

PolyChord (Handley et al., 2015) is a nested sampling algorithm tailored for high-dimensional parameter spaces. What distinguishes different nested sampling algorithms is the method used to sample within the prior space under the given likelihood constraint. PolyChord employs slice sampling at each iteration to sample within an iso-likelihood contour that satisfies the hard likelihood constraint. It is able to identify and evolve separate modes of a posterior semi-independently, which is useful for multi-modal problems.

2.3.1 Slice Sampling

Slice sampling (Neal, 2003) is a method that generates samples from a target distribution proportionally to its density function. It operates by utilising stepping-out and shrinkage procedures to expand the sampling interval, ensuring an unbiased and representative sample.

For a single-variable (one-dimensional) case, the process is as follows: first, a vertical level L is drawn uniformly from zero to the value of the seed point x_0 (i.e. from $[0, \mathcal{L}(x_0)]$). Next, an interval of width w is randomly positioned to encompass x_0 and then expanded in steps until both ends of the interval lie outside the slice. Finally, a new point, x_1, is uniformly sampled from this interval; if a point is sampled outside the slice, the interval is contracted (via shrinkage) until a point inside the slice is found.

2.3.2 Multi-dimensional Slice Sampling

In PolyChord, N-dimensional slice sampling is performed by first applying a linear transformation that converts a degenerate contour into one with dimensions in all directions; this creates a 'whitened' unit hypercube sampling space. A live point x_0 is then randomly chosen in this sampling space. From x_0, a new point x_1 is determined by selecting a random direction and performing one-dimensional slice sampling in that direction. This process is repeated n_{dims} times to generate a new uniformly sampled point x_N that is decorrelated from x_0.

2.3.3 Clustering

In PolyChord, a variant of the k-nearest neighbours algorithm (Cover and Hart, 1967) is adopted for cluster recognition: two points are classified as belonging to the same cluster if they are each other's nearest neighbours. The algorithm iterates k from 2 onwards until the cluster decomposition remains unchanged with the increasing k.

To generate new points accurately, it tracks prior volume estimates as in Equation (2.9), ensuring that new points within a mode are generated with probabilities proportional to the estimated volume. The error in this estimate contributes to the overall evidence error. This proportional allocation ensures that the sampling density accurately reflects the relative sizes of the modes, leading to more precise evidence calculations and reducing the risk of overlooking significant modes during sampling. It effectively treats clusters as separate entities, allowing samples to be generated independently within each mode.

2.4 REACH Data Analysis Pipeline

REACH (de Lera Acedo et al., 2022) is a global 21-cm experiment built to measure the sky-averaged HI 21-cm line based on Bayesian analysis. It distinguishes itself from other similar experiments by utilising physically motivated models of the signal, the foregrounds, and the instrument itself. The REACH data analysis pipeline is developed to perform continuous, physically motivated simulations using parametrised models of the radio foregrounds and the chromatic distortions from the antenna to explain and fit for systematics in a physically understandable way (Anstey et al., 2021). One of the main features of the pipeline is its foreground model, formulated by splitting the sky into spectral index regions. The pipeline includes various global signal models and noise models, the best performing model combination of which can be determined via Bayesian analysis. The nested sampling algorithm PolyChord is implemented in this pipeline.

2.4.1 Foreground Modelling

The first step of simulating an all-sky map is to model the spatial distribution of power around the global 21-cm observation band. In the case of REACH, it is based on the 2008 Global Sky Model (GSM), specifically an instance at 230 MHz. The frequency of 230 MHz is selected for its proximity to the 21-cm observing band while remaining outside it, thereby avoiding contamination by the redshifted 21-cm signal. Radio emissions in the $50-200$ MHz band are predominantly synchrotron and free-free radiation, which means it can be approximated by a power law. Therefore, the scaling is done using the equation

$$T_{\mathrm{fg}}(\Omega,\nu) = (T_{230}(\Omega) - T_\gamma)\left(\frac{\nu}{230}\right)^{-\beta(\Omega)} + T_\gamma \tag{2.13}$$

where β represents the spectral index and T_γ represents the radio background temperature, set at 2.725 K. The spectral index is treated as a function of sky position, $\beta(\Omega)$, because observations at low frequencies have shown it to have significant spatial variations (Guzmán et al., 2011). To account for this variation, a spectral index map is derived by mapping every pixel of a GSM instance at 408 MHz, $T_{408}(\Omega)$, onto the corresponding pixel of a GSM instance at 230 MHz, $T_{230}(\Omega)$:

$$\beta(\Omega) = \frac{\log\left(\frac{T_{230}(\Omega)-T_\gamma}{T_{408}(\Omega)-T_\gamma}\right)}{\log\left(\frac{230}{408}\right)}. \tag{2.14}$$

A sky model, $T_{\mathrm{sky}}(\Omega,\nu)$, is generated across the observing band by dividing the sky into N regions with similar spectral indices and scaling from the base map within each region. A distinct, uniform spectral index is assigned in each region. The required complexity of a foreground model, decided by the value of N, can be derived from the data using Bayesian evidence. The simulated sky map, rotated for the antenna location and observing time, is then convolved with an antenna beam directivity model $B(\Omega,\nu)$ to produce a model of the foregrounds and chromaticity, parametrised by the spectral indices of the regions, at each observing frequency:

$$T_{\mathrm{data}} = \frac{1}{4\pi}\int_0^{4\pi} B(\Omega,\nu)\int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} T_{\mathrm{sky}}(\Omega,\nu,t)\,\mathrm{d}t\,\mathrm{d}\Omega + \hat{\sigma}, \tag{2.15}$$

where $\hat{\sigma}$ represents uncorrelated Gaussian noise. This produces a parametrised model that accounts for the foregrounds and the antenna chromatic distortion.

2.4.2 Bayesian Analysis

A model made up of the foreground model, a signal model, and a noise model is fitted to the input dataset. The REACH data analysis pipeline requires neither the knowledge of a perfectly achromatic antenna nor a foreground known to a high degree of accuracy. However, the parameters of the foregrounds need to be calculated iteratively. This is realised by the nested sampling algorithm PolyChord

(Section 2.3). Apart from the prior range of each parameter, a likelihood function needs to be defined in order to perform the fit.

Likelihood

For a foreground model of only one time-step, a Gaussian likelihood of the form

$$\log \mathcal{L} = \sum_i \left[-\frac{1}{2} \log(2\pi\theta_\sigma) - \frac{1}{2} \left(\frac{T_{\mathrm{D}i} - T_{\mathcal{M}i}(\theta_\mathrm{F}, \theta_\mathrm{S})}{\theta_\sigma} \right)^2 \right], \tag{2.16}$$

is used, where i indexes frequency bins. T_D represents observational data, and $T_\mathcal{M}$ is the parametrised model including both the foreground and the signal model. θ_F, θ_S, and θ_σ are foreground, signal, and noise parameters, respectively. In the case where both the data and foreground model are integrated over time, the likelihood needs to account for all the time bins:

$$\log \mathcal{L} = \sum_i \left[-\frac{1}{2} \log(2\pi\theta_\sigma) - \frac{1}{2} \left(\frac{\frac{1}{N_t} \sum_j \left[T_{\mathrm{D}i,j}\right] - \frac{1}{N_t} \sum_j \left[T_{\mathcal{M}i,j}(\theta_\mathrm{F}, \theta_\mathrm{S})\right]}{\theta_\sigma} \right)^2 \right], \tag{2.17}$$

where j indexes time bins. In the time-dependent local sidereal times (LST) model, datasets at different time-steps are not integrated over the observing time and instead are fitted jointly to their corresponding models using the same parameter samples for all (Anstey et al., 2023). As such, the likelihood of the LST model is formulated as

$$\log \mathcal{L} = \sum_i \sum_j \left[-\frac{1}{2} \log(2\pi\theta_\sigma) - \frac{1}{2} \left(\frac{T_{\mathrm{D}i,j} - T_{\mathcal{M}i,j}(\theta_\mathrm{F}, \theta_\mathrm{S})}{\theta_\sigma} \right)^2 \right]. \tag{2.18}$$

2.4.3 Comparing Different Models

Throughout this thesis, log Bayes ratio $\Delta \log \mathcal{Z}$ is calculated to determine which of the adopted models is favoured. It is calculated by taking the difference between the log of the Bayesian evidence of two competing models, $\mathcal{M}_1$ and $\mathcal{M}_2$:

$$\Delta \log \mathcal{Z} = \log \left(\frac{\mathcal{Z}_{\mathcal{M}_1}}{\mathcal{Z}_{\mathcal{M}_2}} \right) = \log \mathcal{Z}_{\mathcal{M}_1} - \log \mathcal{Z}_{\mathcal{M}_2}. \tag{2.19}$$

Moreover, the significance of a detection is quantified by the log Bayes factor. Throughout this work, this will be referred to as log Bayes ratio and denoted as $\Delta \log \mathcal{Z}$. This value is calculated by subtracting the log evidence of the simpler, foreground-only model, $\log \mathcal{Z}_\mathrm{ns}$, from that of the model which includes a signal, $\log \mathcal{Z}_\mathrm{s}$. The calculation is thus:

$$\Delta \log \mathcal{Z} = \log \mathcal{Z}_\mathrm{s} - \log \mathcal{Z}_\mathrm{ns}. \tag{2.20}$$

A large, positive value indicates strong evidence for a detection, while a large, negative value indicates strong evidence against it.

In order to check whether the fitted signal contains unwanted parts from mistakenly capturing powers contributed by the foreground, the root-mean-square error (RMSE) between the detected signal, T_{21}^{*}, and injected signal, T_{21}, is also calculated:

$$\mathrm{RMSE} = \sqrt{\sum_{i=1}^{n} \frac{(T_{21}(\nu_i) - T_{21}^{*}(\nu_i))^2}{n}}. \tag{2.21}$$

As likelihood and evidence values are yielded by summing over all data points, more time-steps (i.e., more data points) can result in higher overall values.

Part II

Chromatic Effects and Foregrounds

Chapter 3

Ionospheric Effects

When I have seen such interchange of state,
Or state itself confounded to decay...

W. Shakespeare

The research presented in this chapter was originally published in the Monthly Notices of the Royal Astronomical Society as Shen et al. (2022).

3.1 Introduction

As mentioned in Section 1.4, the antennae exhibit chromaticity when operated across large fractional bandwidths. Since neither the chromatic antenna beam nor the sky temperature is smooth as a function of frequency, their convolution disrupts the assumed spectral smoothness of the foreground after integration. The ionosphere introduces extra factors to chromatic mixing by absorption and refraction, and its effects scale approximately as ν^{-2}. Its magnitude can be 2 to 3 orders larger than the global 21-cm signal within the measurement bandwidth. In this chapter, the ionospheric layers that cause significant chromatic distortion in the foregrounds by refraction and absorption are modelled.

This work is based on the paper, Shen et al. (2021), an investigation into how the chromatic ionospheric effects distort the global 21-cm signal and how they behave over time. In this chapter, the challenges posed by time-varying chromatic ionospheric effects on the REACH data analysis pipeline are investigated. Different test cases are compared using log evidence and the signal RMSE.

In this work's formalism, the chromatic ionospheric effects are treated as a distortion of the antenna beam pattern, which is therefore considered to be

time-variable and dependent on the ionospheric conditions. The antenna temperature data subjected to the time-varying chromatic ionospheric effects in the $50 - 200$ MHz band is analysed using the REACH data analysis pipeline (Anstey et al., 2021), in which the antenna beam used to fit the foregrounds also is modified by the simulated ionospheric effects via the two-layer ionospheric model. The foregrounds are modelled using one constant beam for data collected at different times, and then time-varying beams are implemented through the multiple time steps tested in the simulation. The error tolerance of the data analysis pipeline is also examined with respect to the uncertainties in the parameters used to model the ionosphere. To find the error threshold, fixed deviations are introduced to the ionospheric parameters when generating the simulated data.

In this chapter, the REACH data analysis pipeline is employed as a realistic framework to quantify the impact of ionospheric effects on global 21-cm signal extraction. The primary objective of this work is to translate pipeline performance metrics, such as Bayesian evidence and signal RMSE, into practical constraints on observing strategies and data analysis methodologies. To achieve this, the analysis seeks to determine: first, the timescale over which ionospheric variations can corrupt integrated data, and second, the requisite accuracy for ionospheric modelling needed to ensure a credible signal detection. The findings are intended to inform the design of future experiments and establish criteria for the validation of a claimed 21-cm signal.

3.1.1 The Morphology of the Ionosphere

The ionosphere is the ionised portion of the upper atmosphere that envelopes the Earth, extending from approximately 50 km to beyond 1000 km above the Earth's surface; this vast region begins within the mesosphere, extends upwards through the thermosphere, and continues into the exosphere. It is characterised by a significant number of free thermal electrons and ions, generated through the ionisation of the neutral particles by solar radiation and collisions with energetic particles that penetrate the atmosphere. The plasma parameters in Earth's ionosphere exhibit significant variation with location, time, season, solar cycle, and magnetic activity. These variations arise from the coupling, time delays, and feedback mechanisms within the ionosphere-thermosphere system[a]. Additionally, the ionosphere's coupling to other regions in the solar-terrestrial environment, including the Sun, the interplanetary medium, the magnetosphere, and the mesosphere, further complicates its behaviour (Schunk and Nagy, 2000).

The primary source of plasma for the ionosphere is the photoionisation of neutral molecules by solar extreme ultraviolet (EUV) and soft X-ray radiation,

[a]The thermosphere is the neutral atmospheric layer in which the ionospheric plasma is embedded. The two systems are co-located and dynamically coupled.

although other processes, such as ultraviolet (UV) and hard X-ray radiation, may dominate in certain regions. The dominance of these different ionising processes at various altitudes is a consequence of their varying penetration depths: higher-energy radiation from the Sun penetrates deeper into the atmosphere, while lower-energy radiation is absorbed at higher altitudes. Due to its complex refractive index, this magnetised plasma significantly distorts and attenuates incoming radiation as it propagates through (Vedantham et al., 2014).

Although different physical processes dominate at different latitude domains (see Figure 3.1), the variation in electron density with altitude exhibits the same fundamental structure across all latitudes. The ionosphere is traditionally divided into three primary layers, D, E, and F[b], arranged alphabetically with increasing altitude (Chapman, 1931). This layered structure is characterised by variations in electron density and ionic composition, resulting from ionisation by EUV and X-rays from the Sun (Davies, 1990; Evans and Hagfors, 1968; Ratcliffe, 1973). Specifically, the F layer is ionised by EUV radiation, the E layer by soft X-rays and some EUV, and the D layer by hard X-rays and Lyman-α radiation (Davies, 1990).

D layer

The D-layer is the lowest and densest region of the ionosphere, composed primarily of the neutral species, N_2, O_2, and O. Both positive and negative ions also exist, along with some free electrons. The lower part of the D layer is dominated by the cluster ions, which are primarily formed from NO^+ and O_2^+ via hydration (i.e. by attaching to water molecules). The high neutral density leads to a substantial electron collision frequency, causing significant absorption of radio waves.

E layer

In the E layer, the predominant ions are NO^+, O_2^+, and N_2^+. The total ion density is on the order of 10^5 cm^{-3} compared to the neutral density, which exceeds 10^{11} cm^{-3}. As a result, the plasma is weakly ionised in the E layer, and collisions between charged particles are not significant. Therefore, the E layer does not contribute significantly to the 21-cm chromatic distortion.

F layer

The F layer is primarily composed of the atomic species O^+ and O. The peak ion density reaches approximately 10^6 cm^{-3}, while the neutral density is 10^8 cm^{-3}.

[b]Edward Appleton was the first scientist to prove the existence of the ionosphere through systematic radio wave experiments in the 1920s. When he named these layers, he began with the letter E for the one he first discovered, followed by F for the higher layer and D for the lower. The letters A, B, and C were reserved for potential future discoveries that never came (Kivelson and Russell, 1995).

The plasma in the F layer is partially ionised, with collisions occurring both between different charged particles and between charged particles and neutrals. Due to its high electron density, the F layer significantly contributes to the refraction of radio waves.

3.1.2 Refractive Index

The refractive index η of a radio wave travelling through a magneto-ionic plasma such as the ionosphere can be described by the Appleton-Hartree equation (Shkarofsky, 1961). The equation accounts for the effects of electron density, magnetic field, wave frequency, and the direction of wave propagation relative to the magnetic field. The general form of the Appleton-Hartree equation is:

$$\eta^2 = 1 - \frac{X}{1 - iZ - \frac{Y^2 \sin^2\theta}{2(1-X-iZ)} \pm \sqrt{\frac{Y^4 \sin^4\theta}{4(1-X-iZ)^2} + Y^2\cos^2\theta}}. \tag{3.1}$$

The equation's behaviour is governed by three dimensionless parameters. The X term represents the plasma density and is the primary cause of refraction. The Y term represents the magnetic field strength; its effect, which depends on the geometry of the interaction, is to induce birefringence by splitting the wave into two distinct modes. Finally, the Z term is proportional to the electron collision frequency and is responsible for absorption. The mathematical forms of these parameters are:

$$\begin{aligned} X &= \nu_p^2/\nu^2 = (\text{plasma freq.})^2/(\text{radio wave freq.})^2, \\ Y &= \nu_g/\nu \;\; = (\text{electron gyro freq.})/(\text{radio wave freq.}), \\ Z &= \nu_c/\nu \;\; = (\text{electron collision freq.})/(\text{radio wave freq.}), \end{aligned} \tag{3.2}$$

where θ is the angle between the direction of the wave propagation and the magnetic field. The $\pm$ sign differentiates between the ordinary (O wave) and extraordinary (X wave) wave modes[c], which are split due to the magnetic field. The angle θ affects how the magnetic field modifies the propagation, splitting the wave into the O-mode and X-mode due to birefringence.

Since the Earth's magnetic field is assumed to have a negligible effect on the refractive index in this model (Vedantham et al., 2014), the magnetic field term can be omitted to simplify the calculation, by setting electron gyrofrequency and thus Y to zero. The resulting simplified Appleton-Hartree equation, excluding the magnetic field term, is:

$$\eta^2 = 1 - \frac{(\nu_\mathrm{p}/\nu)^2}{1 - i(\nu_\mathrm{c}/\nu)}. \tag{3.3}$$

[c] In a birefringent material like the ionosphere, light splits into two components: the ordinary wave mode, which travels at a constant speed, and the extraordinary wave mode, whose speed depends on its direction.

The electron plasma frequency can be determined by electron density N_e (Rybicki and Lightman, 1986) via the equation:

$$\nu_\mathrm{p} = \frac{1}{2\pi}\sqrt{\frac{N_\mathrm{e}e^2}{m_\mathrm{e}\epsilon_0}}, \tag{3.4}$$

where e is the electron charge, m_e is the effective mass of the electron, and ϵ_0 is the permittivity of free space. The electric field of a plane wave travelling through a homogeneous ionospheric layer is described by

$$E(\Delta s) = E_0 \exp\left(-\frac{i2\pi\nu\Delta s}{c}\eta\right), \tag{3.5}$$

where c is the speed of light, Δs is the distance along the propagation path, and E_0 is the initial electric field. The real part of the refractive index primarily depends on electron density and introduces a change in the phase velocity, resulting in refraction. On the other hand, the imaginary part is mainly influenced by the electron collision rate, which depends on electron density, gas density, and temperature. This imaginary component causes an exponential attenuation of the wave amplitude, leading to absorption.

This ionospheric model consists of two components: one is the F layer, accounting for refraction, characterised by low atmospheric gas density and high electron density, and the other is the D layer, accounting for absorption, with high atmospheric gas density and low electron density. Collisions are negligible in the F layer, so it is modelled using only the real and frequency-dependent term η_F,

$$\eta_F = \left(1 - (\nu_\mathrm{p}/\nu)^2\right)^{1/2}, \quad \nu_\mathrm{c} = 0. \tag{3.6}$$

In contrast, the D layer is modelled with the imaginary and frequency-dependent term η_D, which accounts for absorption due to collisions:

$$\eta_D \approx -\frac{1}{2}\frac{\nu_\mathrm{p}^2\nu_\mathrm{c}/\nu}{\nu^2 + \nu_\mathrm{c}^2}. \tag{3.7}$$

3.2 Modelling Chromatic Ionospheric Effects

Ionospheric refraction is caused by free electrons, released by ionised molecules, that oscillate in response to the electric field of incoming radio waves. The F-layer's high density of these electrons is the main factor affecting the ionosphere's refractive index. The refractive index of a plasma, such as the ionosphere, depends on the electron density and the frequency of the incoming radio wave. As incoming radio waves pass through the ionosphere, the varying electron density with altitude causes a gradual bending of the wave path.

On the other hand, ionospheric absorption is caused by high electron collision frequency. Electron collision frequency represents how often the free electrons in the ionosphere collide with the neutral atoms and molecules in the atmosphere.

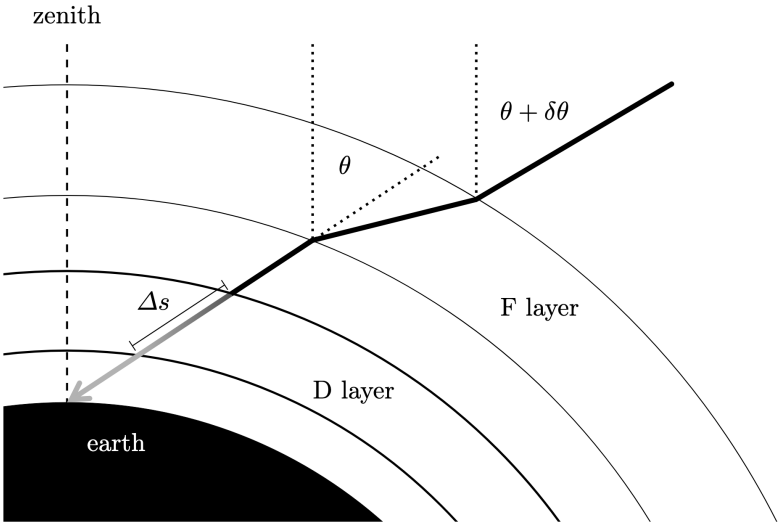


Figure 3.1: Not-to-scale diagram of ionospheric refraction (F layer) and absorption (D layer) of homogeneous ionospheric layers. The dashed line indicates a ray path from the zenith. An incoming ray, shown by the thick solid arrow, first passes through the F layer, where the effective zenith angle decreases by $\delta\theta$ due to F-layer refraction. The incoming ray then reaches the D layer, where the intensity progressively decreases due to D-layer absorption. Δs is the distance along the propagation path in the D layer.

These collisions transfer energy from the radio wave to the ionospheric particles, causing the wave to lose energy. Absorption is significant in the D layer due to high electron collision frequency, particularly for low-frequency radio waves. Figure 3.1 is a not-to-scale diagram of these ionospheric effects.

3.2.1 Effective Beam with Ionospheric Effects

In this model, ionospheric effects are incorporated into the antenna beam model rather than the sky model for analytical convenience. The previously introduced refractive index is used to simulate the effective beam.

Due to Earth's curvature, an incoming ray deviates by $\delta\theta$ as it passes through the F-layer boundaries. This deviation angle can be calculated by integrating the effects of Snell's law incrementally across the layer, where the refractive index at each point is given by Equation (3.3). In a simplified homogeneous F-layer model, where the refractive index is uniform, the complex path integral is unnecessary. Instead, the total deviation angle can be found by simply applying Snell's law at the upper and lower boundaries where the electron density, N_e, and thus the refractive index, changes abruptly.

The F-layer refraction results in different weightings for sky regions with varying deviation angles. Radiation arriving from near the horizon may be refracted to a new zenith angle less than 90°. As a result, the antenna temperature is expected

Table 3.1: List of different types of beams.

	Beam type
B	original antenna beam
$\hat{B}$	antenna beam including ionospheric effects
$\hat{B}_t$	antenna beam including time-dependent ionospheric effects

to increase by approximately $\sim 1-10$ K. The effective antenna beam, accounting for F-layer refraction, can be expressed as:

$$\hat{B}(\nu,\theta,\phi) = B(\nu,\theta+\delta\theta(t),\phi). \tag{3.8}$$

The model assumes a homogeneous D-layer shell. The amplitude of a plane wave travelling in the D layer is given by:

$$|E(\Delta s)| = E_0 \exp\left(\frac{2\pi\nu\Delta s}{c}\eta_D\right). \tag{3.9}$$

Since the intensity of an electromagnetic wave is proportional to the square of its amplitude, the loss factor is

$$\Lambda(\nu,\theta) = \exp\left(\frac{4\pi\nu\Delta s}{c}\eta_D\right). \tag{3.10}$$

The effective beam, subject to absorption in the D layer, can then be written as

$$\hat{B}(\nu,\theta,\phi) = B(\nu,\theta,\phi)\Lambda(\nu,\theta). \tag{3.11}$$

The total propagation path can be integrated numerically using Snell's law. This integral can be approximated by

$$\Delta s \approx \Delta h_D \left(1+\frac{h_D}{a}\right)\left(\cos^2\theta + \frac{2h_D}{a}\right)^{-\frac{1}{2}}, \tag{3.12}$$

where Δh_D is the width of the D layer, h_D is the mean D-layer height, and a is Earth's radius.

Combining Equation (3.8) and (3.11), the effective beam accounting for both the F-layer refraction and the D-layer absorption can be written:

$$\hat{B}(\nu,\theta,\phi) = B(\nu,\theta+\delta\theta(t),\phi)\Lambda(\nu,\theta+\delta\theta(t)). \tag{3.13}$$

Table 3.1 lists the different types of beams mentioned in this thesis.

3.2.2 Time-varying Ionospheric Effects

Ionospheric conditions are time-variable, and the resulting chromatic effects of absorption and refraction do not average out over time. The model used in this work is therefore informed by observational data from the Lowell GIRO Data Center at station Louisvale, South Africa[d]. However, direct nighttime measurements for

[d]The data can be found at https://giro.uml.edu/didbase/.

the D-layer are often unavailable due to its low electron density. To address this, the D-layer parameters in this model are scaled from the available F-layer data based on the relationships established in a fiducial case from a previous study by Vedantham et al. (2014).

When the ionosphere is time-varying, the effective beam must also be treated as time-dependent. This is incorporated into the model by modifying Equation (3.13) as follows:

$$\hat{B}_t(\nu, \theta, \phi, t) = B(\nu, \theta + \delta\theta(t), \phi, t)\Lambda(\nu, \theta + \delta\theta(t), t). \tag{3.14}$$

As described in Section 2.4.1, a sky model $T_{\rm f}(\nu, \theta, \phi)$ is generated across the observing band by dividing the sky into N regions of similar spectral indices and scaling a base map assuming a distinct uniform spectral index in each region. The number of regions is set to be 25 in this work. This sky map is then convolved with an antenna beam model $B(\nu, \theta, \phi)$ to produce a model of foregrounds and chromaticity parametrised by the spectral indices of the regions.

The foregrounds are modelled using two different types of effective beams to account for ionospheric effects. In the first case, a single, static effective beam $\hat{B}$ is used for all time steps. This assumes a constant ionospheric effect over the observation, allowing the time integral in Equation (2.15) to be performed on the sky model alone before the spatial convolution:

$$T_{\rm data} = \frac{1}{4\pi} \int_0^{4\pi} \hat{B}(\nu, \theta, \phi) \int_{t_{\rm start}}^{t_{\rm end}} T_{\rm f}(\nu, \theta, \phi)\, {\rm d}t\, {\rm d}\Omega + \hat{\sigma}, \tag{3.15}$$

In the second case, a time-varying effective beam, $\hat{B}_t$, is used. This requires a joint integral over both time and solid angle:

$$T_{\rm data} = \frac{1}{4\pi} \int_0^{4\pi} \int_{t_{\rm start}}^{t_{\rm end}} \hat{B}_t(\nu, \theta, \phi, t) T_{\rm f}(\nu, \theta, \phi)\, {\rm d}t\, {\rm d}\Omega + \hat{\sigma}. \tag{3.16}$$

In this work, the foreground model is fitted using either a constant beam across all time-steps or multiple beams that correspond to the different effective beams at various time-steps. This approach helps estimate how frequently the beam should be recalculated to achieve meaningful detection. Fixed errors are introduced in the ionospheric parameters when generating the antenna temperature data; specifically, a fixed percentage of increment in electron density, $N_{\rm e}$, collision frequency, $\nu_{\rm c}$, and layer thickness. The resulting data are then fitted using the original, unperturbed beam in the foreground model. This approach allows the error tolerance of the foreground model to be quantified with respect to inaccuracies in the ionospheric parameters.

3.3 Results

In practice, data are collected chronologically at points with a fixed interval between each during observation time and then analysed. In this work, however,

there are gaps between the selected time steps in most of the cases. This is because the ionosphere can suffer significant changes on timescales as short as a few minutes to an hour. One of the aims of this work is to quantify the error tolerance of the REACH data analysis pipeline; therefore, the analysis begins with the more ideal cases. With this in mind, the time steps used for the error analysis were not selected chronologically. Instead, they were chosen as a set of discrete moments known to yield better signal detections when analysed jointly. Consequently, the specific time intervals between these chosen data points are not a relevant variable in this study. The injected Gaussian signal adopted in this work has an amplitude of 0.53 K and a standard deviation of 18.7 MHz centred at 78.1 MHz, which is similar to the absorption feature reported by EDGES (Bowman et al., 2018).

To determine how frequently the beam model should be updated to account for a changing ionosphere, two foreground fitting methods are compared. The first method fits the data from all time steps using a single, constant beam model. The second takes a set of time-varying beam models, where each beam is distorted by the ionospheric conditions corresponding to different time-steps in the observation. The simulated data for these tests are generated for an antenna at the Karoo radio reserve in South Africa (latitude: 30.7°S, longitude: 21.4°E, height: 25 m).

3.3.1 Foregrounds Modelled with a Constant Beam

A primary method for enhancing sensitivity in radio astronomy is temporal integration. This approach, however, presumes instrumental and environmental stability over the integration period. This assumption is tested by combining data from multiple time steps, which are then modelled using a single, static beam model representative of the observation.

Two time steps

Simulated antenna temperature data are selected at four different time steps, $[t_1, t_2, t_3, t_4]$, based on their better performance when analysed in pairs with the reference time step, t_0. The selection criteria required that each pair, when modelled with a single beam, B_{t_0}, yield a log Bayes ratio of $\Delta \log \mathcal{Z} > 10$ and an RMSE < 0.1 K using the LST version of the REACH data analysis pipeline. These specific time steps, along with their corresponding evidence values, are listed in Table 3.2 and form the basis for the subsequent results..

Since the model settings are identical across the four two-time-step cases, their log evidence values ($\log \mathcal{Z}_s$) are comparable. A higher $\log \mathcal{Z}_s$ could be a sign of better signal detection. The fitted signal and foreground residual plots of case $\mathrm{C}_1[t_0, t_1]$ are shown in Figure 3.2. In these cases, the signal, shown in blue contours, is detected with high significance.

Table 3.2: Log Evidence and Signal RMSE for the explored cases (C_i). The **top panel** ($C_1 - C_9$) shows results where the foregrounds are modelled using one single beam (B_{t_0}). The **bottom panel** ($C_{10} - C_{11}$) shows those with the foregrounds modelled using two beams; the foregrounds of the time step coloured in red are modelled using its corresponding beam, and the rest using B_{t_0}. The cases coloured in gray are the reference cases from the top panel. $[t_i]$ lists all the time steps analysed in each case.

Case	[t]	RMSE (K)	$\log \mathcal{Z}_s$	$\Delta \log \mathcal{Z}$
one beam				
C_1	t_0, t_1	0.066	231.5 ± 0.3	62.8
C_2	t_0, t_2	0.035	520.2 ± 0.3	180.4
C_3	t_0, t_3	0.056	93.3 ± 0.4	161.8
C_4	t_0, t_4	0.024	519.0 ± 0.3	253.5
C_5	t_0, t_1, t_2	0.082	-4720.2 ± 0.4	194.4
C_6	t_0, t_2, t_4	0.034	752.6 ± 0.4	316.1
C_7	t_0, t_1, t_2, t_4	0.055	-4665.1 ± 0.4	196.5
C_8	t_0, t_2, t_3, t_4	0.083	-7308.7 ± 0.4	53.0
C_9	t_0, t_1, t_2, t_3, t_4	0.063	-9387.3 ± 0.4	103.0
two beams				
C_{10}	t_0, t_1, t_2	0.056	-177.2 ± 0.3	9478.2
C_5	t_0, t_1, t_2	0.082	-4720.2 ± 0.4	194.4
C_{11}	t_0, t_2, t_4	0.022	796.9 ± 0.4	385.5
C_6	t_0, t_2, t_4	0.034	752.6 ± 0.4	316.1

$[t_0, t_1, t_2, t_3, t_4] = [21\text{h}\,00\text{m}, 21\text{h}\,15\text{m}, 45\text{h}\,15\text{m}, 46\text{h}\,15\text{m}, 47\text{h}\,45\text{m}]$

Three or more time steps

Table 3.2 includes the evidence of the cases consisting of three or more time steps; it reveals an insight for this type of joint analysis: for a case with multiple time steps to yield a high evidence value, all of its constituent two-time-step subgroups must also be individually well-fitted. For instance, the three-time-step case $[t_0, t_2, t_4]$ is successful because the pairs $[t_0, t_2]$ and $[t_0, t_4]$ are themselves high-evidence combinations (as shown in Table 3.2). For the other cases, one observes that the evidences drop to negative values, but their respective $\Delta \log \mathcal{Z}$ and RMSE still imply significant detection; Figure 3.3 shows the plot of the fitted signals as well as the foreground residuals of case $C_5[t_0, t_1, t_2]$ at each time step, and the reconstructed signals in both time steps are close to the injected signal, as implied by the RMSE. The plots of case $C_6[t_0, t_2, t_4]$ are shown in Figure 3.4. Despite the negative, low evidences for multiple-time-step cases, their RMSEs and reconstructed signals still show significant detection, even in the five-time-step case. The reconstructed signal and foreground residual plots are shown in Figure 3.5. One should note that the time steps selected for the analyses are the more ideal cases, which might not be the case in actual observations.

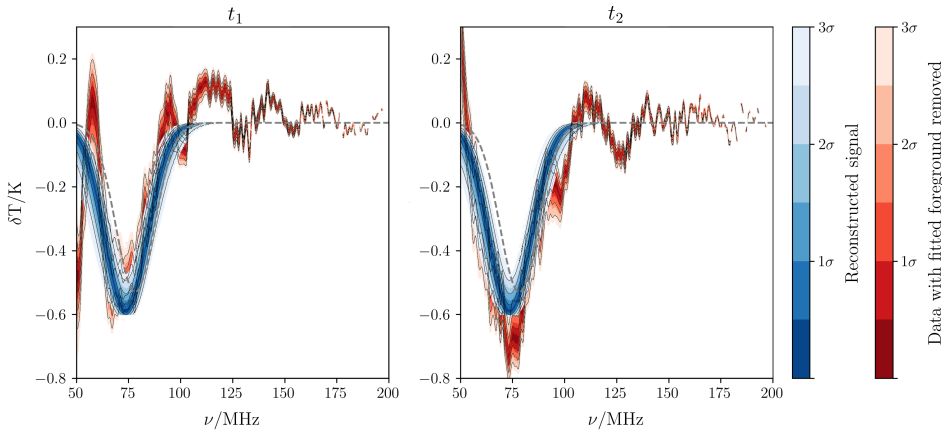


Figure 3.2: A two-time-step case, where the foregrounds are fitted using one single beam B_{t_0} (case $C_1[t_0, t_1]$ of Table 3.2, the signal RMSE is 0.066 K and $\Delta \log \mathcal{Z} = 62$). The reconstructed signal is shown in blue, and the data with the fitted foregrounds removed is shown in red. The gray dashed line is the injected signal. The low residuals, low signal RMSE and high $\Delta \log \mathcal{Z}$ indicate a significant detection.

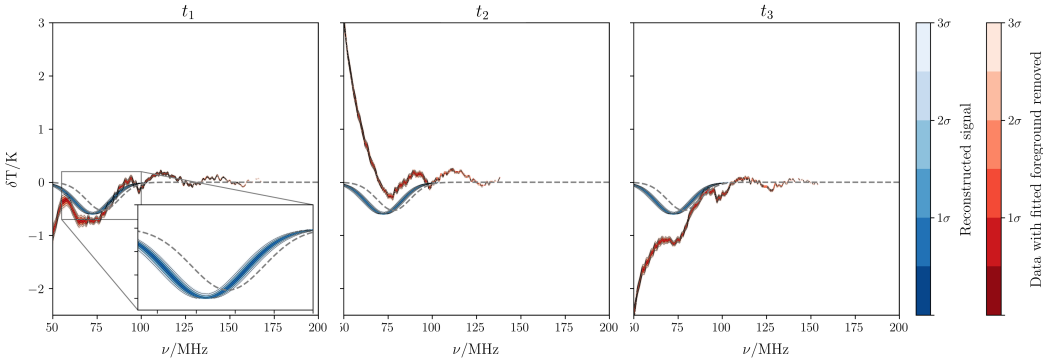


Figure 3.3: A three-time-step case, where the foregrounds are fitted using one single beam B_{t_0} (Case $C_5[t_0, t_1, t_2]$ from Table 3.2, the RMSE of the reconstructed signal is 0.082 K and $\Delta \log \mathcal{Z} = 190$). The reconstructed signal is shown in blue, and the data with the fitted foreground removed is shown in red. The gray dashed line is the injected signal. The signal is recovered with high significance despite the low $\log \mathcal{Z}$ and large foreground residuals in the last two time steps.

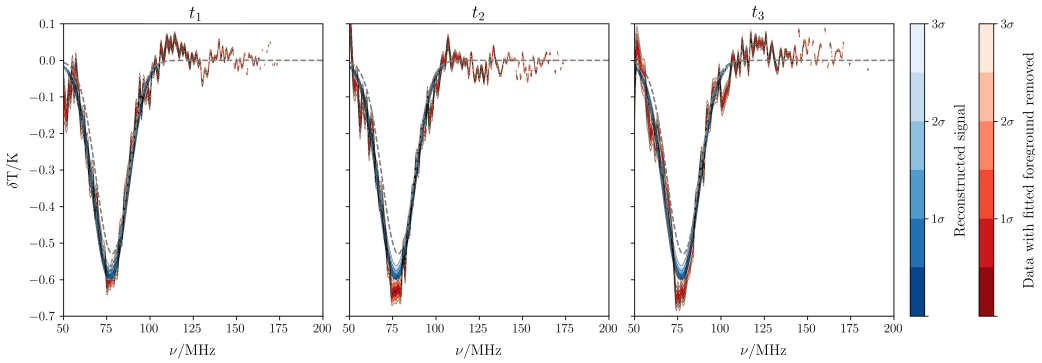


Figure 3.4: A three-time-step case, where the foregrounds are fitted using one single beam B_{t_0} (Case $C_6[t_0, t_2, t_4]$ from Table 3.2, the RMSE of the reconstructed signal is 0.034 K and $\Delta \log \mathcal{Z} = 317$). The reconstructed signal is shown in blue, and the data with the fitted foreground removed is shown in red. The gray dashed line is the injected signal. The signal is recovered with high significance, and the foreground residuals are small across all time steps.

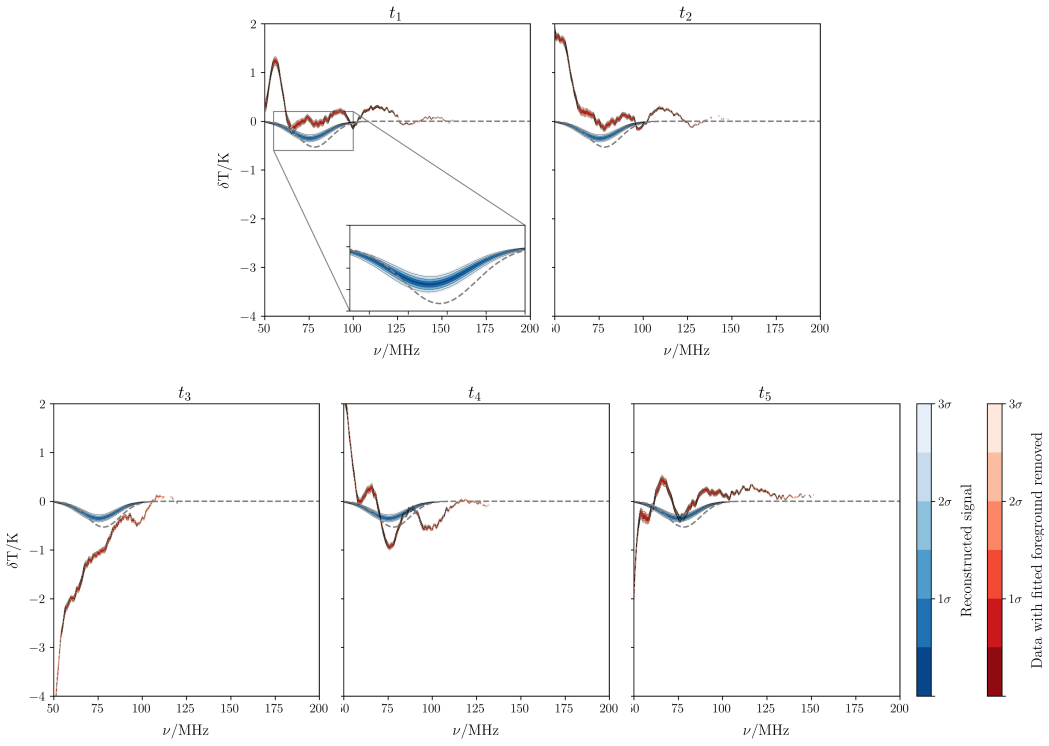


Figure 3.5: A five-time-step case, where the foregrounds are fitted using one single beam B_{t_0} (Case $C_9[t_0, t_1, t_2, t_3, t_4]$ from Table 3.2, the RMSE of the reconstructed signal is 0.063 K and $\Delta \log \mathcal{Z} = 103$). The reconstructed signal is shown in blue, and the data with the fitted foreground removed is shown in red. The gray dashed line is the injected signal. The signal is recovered with high significance despite the low evidence $\log \mathcal{Z}$ and large foreground residuals.

3.3.2 Foregrounds Modelled using Time-varying Beams

The results shown in Section 3.3.1 where the foreground signal is modelled using only one beam for multiple time steps demonstrate that this approach can be problematic in the presence of the chromatic effects introduced by a changing ionosphere. Therefore, in this section, time-varying beams are tested to see how much the analyses can be improved by introducing more corresponding beams when fitting the foregrounds for cases with multiple time steps.

This section tests how an analysis is affected by adding an extra time step. Specifically, it explores how the signal detection for a two-time-step case, where the foreground is modelled with a single constant beam, differs when a third time step, modelled with its own corresponding beam, is introduced. The bottom panel of Table 3.2 shows two such cases. This extra time step is t_2, shown in Table 3.2 in red. One observes that although $[t_0, t_1]$ as a pair does yield good evidence, the evidence when analysed together with another time step t_2 (case C_{10}), modelled using its corresponding beam, drops to a negative value. The other case C_{11}, on the other hand, as $[t_0, t_4]$ yields good evidence as a pair, yields high evidence as well as $\Delta \log \mathcal{Z}$ when analysed with the same time step (t_2) as in case C_{10}. The RMSE also decreases, suggesting better signal detection. This again shows that signal detection could be improved by including more time steps but only if, in smaller subgroups, yield better evidence and signal detection.

3.3.3 Deviation in Modelled Ionospheric Parameters

The effectiveness of a time-varying beam model depends on the accuracy of the ionospheric parameters used in its construction. To quantify the required precision, controlled errors are introduced into the ionospheric parameters of the simulated data: a fixed amount of error [2%, 5%, 10%, 15%] is applied to the ionospheric parameters (a fixed percentage of increment in electron density, N_e, collision frequency, ν_c, and layer thickness), when generating the antenna temperature data to determine how much error would lead to unsuccessful detection. A steady decrease in antenna temperature T_A, due to the applied error increasing the chromatic ionospheric absorption, is shown in Figure 3.6. The data analysis pipeline then fits the foregrounds to the antenna temperature data using the beam distorted by the ionosphere but with no error applied to the parameters.

Two time steps, foregrounds modelled using one beam and two corresponding beams

The results of introducing fixed deviations to the ionospheric parameters are shown in Table 3.3 and Figure 3.7.

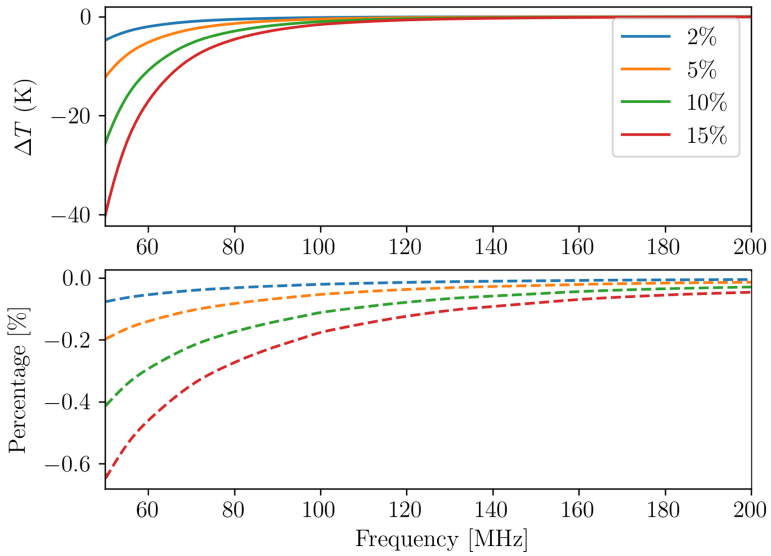


Figure 3.6: The change in integrated antenna temperature as a function of frequency. The different curves correspond to error levels of 2%, 5%, 10%, and 15% applied to the ionospheric parameters, electron density N_e, collision frequency ν_c, and layer thickness. The plot shows both the absolute change in temperature, $\Delta T = T' - T$, and the percentage change, $\Delta T/T$.

As expected, the quality of the signal recovery degrades as the error in the ionospheric model increases. This is shown clearly in the top-left panel of Figure 3.7, where the signal RMSE grows consistently with the applied deviation to the ionospheric parameters, when the foregrounds of all the time steps are modelled using only a single beam B_{t_0}.

$\Delta \log \mathcal{Z}$ generally decreases with this deviation, with the exception of an anomalous increase in the 15% case. This may seem counter-intuitive, but it does not indicate a better signal detection. When the deviation reaches 10%, the RMSE becomes really high and $\log \mathcal{Z}_s$ turns negative.

The likely explanation for the anomaly is that the signal model is no longer fitting the injected 21-cm signal; instead, it is erroneously fitting a feature in the complex foreground residuals. Therefore, the high $\log \mathcal{Z}_s$ does not indicate a successful signal detection.

As the configuration of the analysis pipeline is identical for all the cases with deviations, their $\log \mathcal{Z}_s$ values are directly comparable. The results show that $\log \mathcal{Z}_s$ drops consistently as the error in the ionospheric parameters increases. The value turns negative as the deviation reaches 10%, indicating that a reliable signal detection is no longer viable beyond this threshold.

As expected, using a time-varying beam model improves signal detection compared to using a single, constant beam. The bottom panel of Table 3.3 shows the log evidence for cases with fixed deviations in their ionospheric parameters. In

Table 3.3: Log evidence & RMSE of the reconstructed signal. The cases in this table all consist of two time steps $[\boldsymbol{t_0}, \boldsymbol{t_1}]$, with deviation applied to the ionospheric parameters. The **top panel** shows cases where the foregrounds are modelled using one single beam (B_{t_0}). For those listed in the **bottom panel**, the foregrounds are modelled using two corresponding beams. RMSE is the root-mean-square error of the reconstructed signal with respect to the injected signal. $\log \mathcal{Z}_\mathrm{s}$ is the evidence yielded by the foreground model including a Gaussian signal model, and $\Delta \log \mathcal{Z}$ is the difference between the evidence yielded by the foreground model including a Gaussian signal model and the one without (see Equation 2.20).

Deviation	RMSE (K)	$\log \mathcal{Z}_\mathrm{s}$	$\Delta \log \mathcal{Z}$
one beam			
0%	0.066	231.5 ± 0.3	62.8
2%	0.076	189.1 ± 0.3	42.1
5%	0.091	81.0 ± 0.3	16.6
10%	0.260	-377.1 ± 0.4	0.2
15%	0.283	-1822.5 ± 0.4	2.0×10^{10}
two beams			
0%	0.017	580.2 ± 0.4	208.7
2%	0.071	455.7 ± 0.4	98.1
5%	0.169	164.4 ± 0.3	2.4
10%	0.208	-740.2 ± 0.4	34.7
15%	0.208	-3113.0 ± 0.4	99.6

all scenarios, modelling the foregrounds with two corresponding beams yields a significant improvement over the single-beam approach.

Three time steps, foregrounds modelled using one beam and three corresponding beams

The bottom panel of Table 3.4 shows the results for the three-time-step cases, where the foreground is modelled using three corresponding beams. For cases with low levels of ionospheric deviation, this time-varying beam model yields a significant improvement in the log evidence compared to the single-beam approach.

Similar to the two-time-step results shown in Section 3.3.3, both the single-beam and three-beam models show a reliable detection, but only until the deviation in the ionospheric parameters reaches 10%. The lower panels of Figure 3.7 illustrate this trend, plotting the signal RMSE and $\Delta \log \mathcal{Z}$ as a function of the applied deviation.

3.4 Conclusions

A quantification of the impact of time-varying ionospheric effects on the detection of the global 21-cm signal is presented in this work. The findings indicate that the ionosphere introduces significant temporal variability into the effective instrument

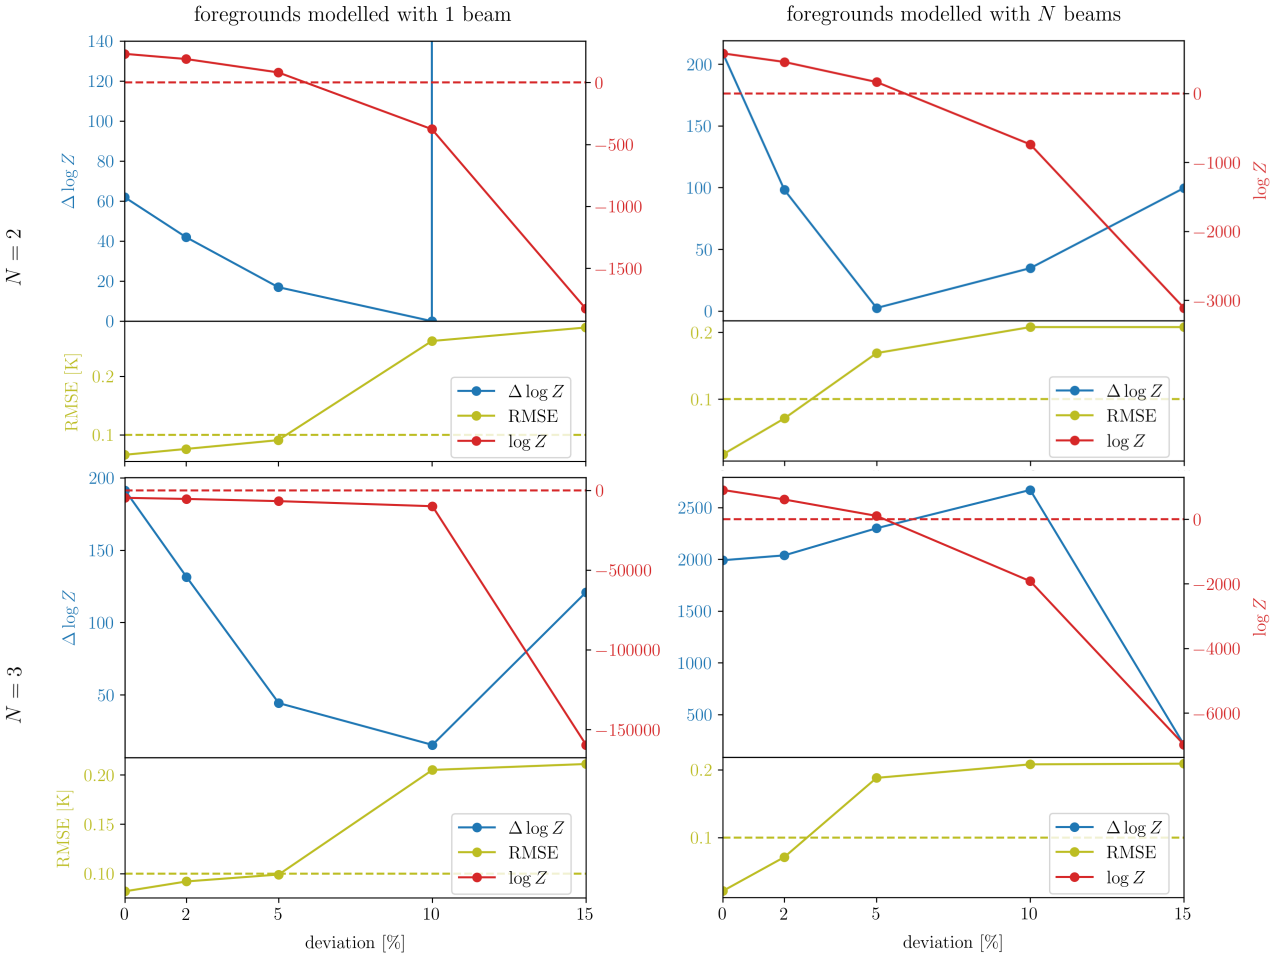


Figure 3.7: The impact of fixed percentage errors in the ionospheric model on. $\Delta \log \mathcal{Z}$ (blue), $\log \mathcal{Z}$ (red), and RMSE (olive). Each panel shows these metrics as a function of the applied deviation at [2%, 5%, 10%, 15%] to the ionospheric parameters, N_e, ν_c, and layer thickness. The foregrounds in the data analysis pipeline are modelled using one single beam (B_{t_0}). The four panels compare different scenarios. **Top-left:** a two-time-step case ($[t_0, t_1]$) modelled with one beam. **Top-right:** a two-time-step case ($[t_0, t_1]$) modelled with two beams. **Bottom-left:** a three-time-step case ($[t_0, t_1, t_2]$) modelled with one beam. **Bottom-right:** a three-time-step case ($[t_0, t_1, t_2]$) modelled with three beams. Reference lines for $\log \mathcal{Z}_s = 0$ (red dashed) and RMSE = 0.10 K (olive dashed) are shown. The results are also shown in Table 3.3 and Table 3.4. In general, there is a consistent increase in RMSE and decrease in $\log \mathcal{Z}_s$. The anomalous increase in $\Delta \log \mathcal{Z}$ at high deviation is due to the model erroneously fitting foreground residuals instead of the injected signal. Overall, it suggests that a reliable detection becomes difficult when the error in the ionospheric model exceeds approximately 5%.

Table 3.4: Log evidence & RMSE of the reconstructed signal. The cases in this table all consist of three time steps $[\boldsymbol{t_0}, \boldsymbol{t_1}, \boldsymbol{t_2}]$, with deviation applied to the ionospheric parameters. For cases listed in the **top panel**, the foregrounds are modelled using one single beam (B_{t_0}). For those listed in the **bottom panel**, the foregrounds are modelled using three corresponding beams. RMSE is the root-mean-square error of the reconstructed signal with respect to the injected signal. $\log \mathcal{Z}_\mathrm{s}$ is the evidence yielded by the foreground model including a Gaussian signal model, and $\Delta \log \mathcal{Z}$ is the difference between the evidence yielded by the foreground model including a Gaussian signal model and the one without (see Equation 2.20).

Deviation	RMSE (K)	$\log \mathcal{Z}_\mathrm{s}$	$\Delta \log \mathcal{Z}$
one beam			
0%	0.082	-4723.2 ± 0.4	191.4
2%	0.092	-5426.3 ± 0.4	131.5
5%	0.099	-6803.2 ± 0.4	44.3
10%	0.205	-9987.9 ± 0.4	15.4
15%	0.211	-15659.0 ± 0.4	121.9
three beams			
0%	0.021	904.4 ± 0.4	1991.5
2%	0.071	609.2 ± 0.4	2038.0
5%	0.188	101.9 ± 0.4	2300.2
10%	0.208	-1919.3 ± 0.4	2669.7
15%	0.209	-6977.1 ± 0.4	210.4

response, rendering standard long-duration integrations with a static instrumental model insufficient.

To study how chromatic ionospheric effects on the foregrounds affect signal detection, cases including multiple time steps are explored using the REACH data analysis pipeline. These cases are then analysed and compared in terms of log evidence and the RMSE of the reconstructed signal. For analytical convenience, the chromatic ionospheric effects are modelled as a distortion of the antenna beam pattern. This creates a time-varying effective beam that changes with ionospheric conditions, even though the physical antenna itself remains static with respect to these specific effects. This modelled variability is purely an ionospheric effect, distinct from any potential intrinsic variations in the physical antenna's response.

To determine the most efficient modelling strategy for signal detection, the analysis explores different approaches to handling the time-varying instrumental response. Apart from using a corresponding, ionosphere-distorted beam for every time bin, the foregrounds are modelled using a reduced set of representative beams. The error tolerance of the pipeline for a successful detection is also investigated with respect to the precision of the ionospheric parameters. The results can be summarised by the following points:

(i) Signal detection does not necessarily improve when more time steps are analysed together using the foreground model fitted with only one beam.

(ii) When the foregrounds are modelled using a reduced set of representative

beams (i.e., fewer beams than time steps), the outcome depends on the compatibility of the data being combined. Strong signal detections, indicated by the low RMSE and high evidence, are achieved when all the constituent two-time-step subgroups are individually well-fitted. While it may still be possible to achieve a statistically significant detection combining less compatible subgroups, the signal recovery is typically degraded.

(iii) Using more corresponding beams at different time steps when modelling the foregrounds significantly improves signal detection.

(iv) Less than 5% error in our knowledge of the ionospheric parameters can lead to unsuccessful detection of the injected global 21-cm signal.

This work establishes that the global signal can be robustly recovered if the uncertainty in ionospheric modelling is kept below 5% to prevent foreground subtraction residuals from compromising the measurement. In multiple-time-step cases where the ionospheric condition is relatively stable, and when our knowledge of the ionospheric condition is sufficient, robust signal recovery can be achieved. However, if the ionosphere varies significantly during an observation, a static beam model is insufficient, and a time-varying beam configuration becomes essential for a statistically meaningful detection.

A successful detection in such cases therefore requires a dynamic approach, incorporating contemporaneous, high-cadence monitoring to inform a time-varying beam model within the analysis pipeline. The established 5% tolerance provides a quantitative framework for designing observing strategies and for verifying future claimed detections of the global 21-cm signal.

Chapter 4

Foreground Polarisation

I am not what I am.

W. Shakespeare

The work presented in this chapter was published in the Monthly Notices of the Royal Astronomical Society as Shen et al. (2025).

4.1 Introduction

Polarised Galactic emissions are Faraday rotated when they pass through the magnetised interstellar medium (Rybicki and Lightman, 1986). The magnitude of this rotation is proportional to the wavelength squared (λ^2), introducing a chromatic effect where the amount of contamination leaking into the total intensity measurement varies with frequency.

However, our direct observational knowledge of this low-frequency polarised emission is sparse. While several pioneering surveys have provided crucial first glimpses, they are often limited in sky coverage or frequency range. For example, interferometers like the Murchison Widefield Array (MWA) and LOFAR Two-meter Sky Survey (LOTSS) have produced deep but spatially limited maps below 200 MHz (Bernardi et al., 2013; Van Eck et al., 2019). Conversely, surveys like the S-band Polarisation All-Sky Survey (S-PASS) have achieved wide sky coverage but at much higher frequencies (2.3 GHz), where Faraday rotation effects are weaker (Carretti et al., 2019). This observational gap is expected to be remedied within the next decade, with mapping diffuse polarisation listed as a primary science goal for the SKA-Low (Chrysostomou et al., 2020).

Contamination due to polarised foregrounds for sky-averaged 21-cm experiments has been studied by Spinelli et al. (2018, 2019); they created a model to simulate contamination by polarised foregrounds through simulating all-sky polarised Galactic diffuse emission at low frequencies based on the statistical properties shown in the patchy MWA 2400 degree^2 survey. Faraday rotation measure synthesis was adopted to generate a three-dimensional representation of the linearly polarised sky. This work investigates the effect of polarised foreground in the extraction of the global 21-cm signal feature. The models to simulate a linearly polarised sky include the one used in Spinelli et al. (2018) as well as other methods built to describe the otherwise complex polarised sky by tunable parameters. The robustness of the REACH pipeline is specifically tested against polarised foreground contamination using its improved Bayesian data analysis approach.

In addition to using three distinct models of polarised foregrounds to test the robustness of the REACH data analysis pipeline (Anstey et al., 2021), the properties of polarisation leakage that pose the greatest threat to a 21-cm detection are also characterised. The aim is to determine whether the primary challenge is the overall amplitude of the leakage or its spectral structure. By simulating scenarios ranging from physically-motivated extrapolations of existing surveys to controlled toy models, pipeline performance metrics are translated into a broader understanding of how the spatial distribution of Faraday depth across the sky impacts the spectral character of the contamination. This allows us to assess the conditions under which polarised foregrounds could be mistaken for the global 21-cm signal itself.

4.1.1 Polarised Synchrotron Radiation

The foreground is composed predominantly of diffuse synchrotron radiation. Synchrotron radiation is emitted by relativistic charged particles spiralling around magnetic field lines. In our Galaxy, it primarily originates from cosmic ray electrons produced by supernova explosions as they interact with the Galactic magnetic field.

Polarisation state

The observed polarisation state of synchrotron radiation depends on the angle θ between the magnetic field and the observer's line of sight.

When the magnetic field is perpendicular to the line of sight, $\theta = 90°$, the acceleration vector of the electron oscillates in a plane perpendicular to the magnetic field, producing linearly polarised radiation. When the magnetic field is parallel to the line of sight, $\theta = 0°$, the electron undergoes circular motion about the magnetic field lines, causing its acceleration vector to rotate with time, resulting in circularly polarised radiation.

For intermediate angles, $0° < \theta < 90°$, the radiation is elliptically polarised, with the ellipse aspect ratio given by $\cos\theta$. Highly relativistic electrons typically possess significant velocity components both perpendicular and parallel to the magnetic field lines. As a result, their beamed radiation is usually observed at an oblique angle rather than purely parallel or perpendicular to the line of sight. This leads to synchrotron radiation from a single electron being generally elliptically polarised.

Relativistic beaming

For a non-relativistic particle undergoing gyro-motion in a magnetic field, the emitted radiation forms an approximate dipole pattern with maximum intensity along the electron's direction of motion, $\mathbf{v}$. However, at relativistic speeds, aberration effects significantly reshape this pattern (Rybicki and Lightman, 1986). Due to Lorentz contraction, the emission is strongly beamed forward into a narrow angular region $\pm\sin\phi \approx \pm\phi \approx \pm 1/\gamma$ from the cone of maximum radiation, where γ is the Lorentz factor. An observer detects radiation only when the electron's velocity vector aligns within this narrow cone along the line of sight. The total opening angle of the beamed radiation is $\Delta\alpha \approx 2/\gamma$.

Observed linear polarisation

Now, consider a distant observer viewing relativistic electrons moving through a magnetised field. For electrons spiralling counter-clockwise along a magnetic field line, radiation received from within the inner $1/\gamma$ region of the velocity cone appears left-hand elliptically polarised because the electrons also appear to move counter-clockwise. Conversely, radiation received from the outer $1/\gamma$ region appears right-hand elliptically polarised. Since both regions contribute approximately equally on either side of the line of sight, their elliptical components cancel out, leaving the observed partially polarised emission predominantly linear.

4.1.2 Faraday Rotation

Faraday rotation occurs when the plane of linear polarisation of an electromagnetic wave rotates while propagating through a magnetised plasma. This occurs because the plasma becomes birefringent under the influence of a magnetic field, causing the left and right circularly polarised components of the wave to travel at different speeds. In the interstellar medium, birefringence primarily arises from the magnetised plasma composed of free electrons and ions.

The following mathematical formulation draws in part on the derivation presented in Ferrière et al. (2021), shows how a magnetised plasma becomes birefringent. The analysis begins by defining the physical system. Consider a cold and magnetised plasma characterised by an electron density n_e and a magnetic field $\mathbf{B}$.

The plasma's response is governed by two characteristic frequencies: the electron plasma frequency, ω_e, defined as $\omega_\mathrm{e} = \sqrt{\frac{4\pi n_\mathrm{e} e^2}{m_\mathrm{e}}}$ and the electron gyrofrequency, Ω_e, given by $|\Omega_\mathrm{e}| = \frac{eB}{m_\mathrm{e} c}$. Here, m_e denotes the mass of the electron, e the electric charge, B the magnetic field strength, and c the speed of light. Now, consider an electromagnetic wave with angular frequency, ω, corresponding to frequency $\nu = \omega/2\pi$ and wave vector is $\mathbf{k}$, where $k = |\mathbf{k}|$ is the wavenumber and $\hat{\mathbf{k}} = \mathbf{k}/k$ is the unit propagation direction. For simplicity, start with a simple case where the wave propagates parallel to the magnetic field, $\mathbf{k} \parallel \mathbf{B}$. Under the approximation that the wave frequency is much higher than the gyrofrequency, $|\Omega_\mathrm{e}| \ll \omega$, the dispersion relation is

$$\begin{aligned} \omega^2 &= c^2 k^2 + \frac{\omega_\mathrm{e}^2}{1 \mp \frac{|\Omega_\mathrm{e}|}{\omega}} \\ &\approx c^2 k^2 + \omega_\mathrm{e}^2 \left(1 \pm \frac{|\Omega_\mathrm{e}|}{\omega}\right). \end{aligned} \tag{4.1}$$

In this expression, the c^2k^2 term represents the electromagnetic wave propagation in free space, while the additional term accounts for the plasma's diamagnetic effect where the induced electric current is primarily due to the dominant electric force, with a much lesser contribution from the magnetic force. The $\mp$ (and $\pm$ in the approximated form) sign corresponds to the two circularly polarised modes. The upper sign ($-$ in the original form and $+$ in the approximated form) represents the right circularly polarised mode, where the electric field vector rotates right-handedly about $\mathbf{B}$; the lower sign represents the left-hand circularly polarised mode, where the electric field vector rotates left-handedly. From this dispersion relation, and assuming the wave frequency is also much higher than the plasma frequency, $\omega_\mathrm{e} \ll \omega$, the phase velocity $v_\phi \equiv \omega/k$ for each mode can be approximated:

$$\begin{aligned} v_\phi \equiv \frac{\omega}{k} &\approx \frac{c}{1 - \frac{\omega_\mathrm{e}^2}{2\omega^2}\left(1 \pm \frac{|\Omega_\mathrm{e}|}{\omega}\right)} \\ &\approx c\left[1 + \frac{\omega_\mathrm{e}^2}{2\omega^2}\left(1 \pm \frac{|\Omega_\mathrm{e}|}{\omega}\right)\right]. \end{aligned} \tag{4.2}$$

This expression demonstrates that the phase velocities for the right and left circularly polarised modes are slightly different, with $v_{\phi,\mathrm{R}} > v_{\phi,\mathrm{L}}$. As these modes propagate away from the source, this velocity difference causes their relative phase to gradually increase. For a wave moving through an inhomogeneous plasma, where the wavenumber k is not constant, the phase must be calculated by integrating along the path:

$$\phi = \omega t - \int_{s_0}^{s_1} k \, \mathrm{d}r, \tag{4.3}$$

assuming the initial phase offset is zero at $r = s_0$.

Rotation of Linearly Polarised Radiation

Consider a source of linearly polarised synchrotron radiation, the initial electric field vector at the source is

$$\mathbf{E}_{\ell *} = E_0 \cos(\omega t)\hat{\mathbf{k}}_*, \tag{4.4}$$

where subscript $*$ indicates that it is at the source. A linearly polarised wave can be understood by decomposing it into left and right circularly polarised modes with the same angular frequency ω. The corresponding electric field vectors can be written likewise:

$$\mathbf{E}_\ell = \mathbf{E}_\mathrm{R} + \mathbf{E}_\mathrm{L}. \tag{4.5}$$

Initially at $r = s_0$, the two modes have the same phase $\phi_* = \omega t$. As they propagate away from the source, the phase difference $\Delta\phi = \phi_\mathrm{R} - \phi_\mathrm{L}$ increases by

$$\begin{aligned} \Delta\phi &= -\int_{s_0}^{s_1} \Delta k \, \mathrm{d}r \\ &= \int_{s_0}^{s_1} \frac{\Delta v_\phi}{v_\phi} k \, \mathrm{d}r \\ &= \int_{s_0}^{s_1} \frac{\omega_\mathrm{e}^2 |\Omega_\mathrm{e}|}{\omega^3} k \, \mathrm{d}r. \end{aligned} \tag{4.6}$$

The sum of the right and left circularly polarised modes remains linear; the change is manifested by the rotated polarisation angle, which is Faraday rotation.

Faraday rotation angle

The initial electric field of the right circularly polarised mode can be described by

$$\mathbf{E}_{\mathrm{R}*} = \frac{1}{2} E_0 \left[\cos(\omega t)\hat{\mathbf{k}}_x + \sin(\omega t)\hat{\mathbf{k}}_y \right], \tag{4.7}$$

and the left circularly polarised mode would then be

$$\mathbf{E}_{\mathrm{L}*} = \frac{1}{2} E_0 \left[\cos(\omega t)\hat{\mathbf{k}}_x - \sin(\omega t)\hat{\mathbf{k}}_y \right]. \tag{4.8}$$

As they propagate away from the source, t would be replaced with the retarded time, and given $v_\phi \equiv \omega/k$ and Equation (4.3):

$$\omega \left(t - \int_{s_0}^{s_1} \frac{\mathrm{d}r}{v_\phi} \right) = \omega t - \int_{s_0}^{s_1} k \, \mathrm{d}r = \phi. \tag{4.9}$$

Thus, the modified field expressions become

$$\mathbf{E}_\mathrm{R} = \frac{1}{2} E_0 \left[\cos\phi_\mathrm{R}\hat{\mathbf{k}}_x + \sin\phi_\mathrm{R}\hat{\mathbf{k}}_y \right], \tag{4.10}$$

$$\mathbf{E}_\mathrm{L} = \frac{1}{2} E_0 \left[\cos\phi_\mathrm{L}\hat{\mathbf{k}}_x - \sin\phi_\mathrm{L}\hat{\mathbf{k}}_y \right]. \tag{4.11}$$

Summing these components gives the resultant electric field, and by using the sum-to-product identities:

$$\begin{aligned}\mathbf{E}_\ell &= \mathbf{E}_\mathrm{R} + \mathbf{E}_\mathrm{L} \\ &= \frac{1}{2}E_0\left[(\cos\phi_\mathrm{R} + \cos\phi_\mathrm{L})\,\hat{\mathbf{k}}_x + (\sin\phi_\mathrm{R} - \sin\phi_\mathrm{L})\,\hat{\mathbf{k}}_y\right] \\ &= E_0\cos\left(\frac{\phi_\mathrm{R}+\phi_\mathrm{L}}{2}\right)\left[\cos\left(\frac{\phi_\mathrm{R}-\phi_\mathrm{L}}{2}\right)\hat{\mathbf{k}}_x + \sin\left(\frac{\phi_\mathrm{R}-\phi_\mathrm{L}}{2}\right)\hat{\mathbf{k}}_y\right].\end{aligned} \tag{4.12}$$

This result reveals that the polarisation angle has rotated by $\Delta\phi/2 = (\phi_\mathrm{R} - \phi_\mathrm{L})/2$, which is the Faraday rotation angle, $\Delta\psi$. The difference in phase velocities between the right and left circularly polarised modes leads to this cumulative phase shift, which manifests as a rotation of the polarisation plane.

Rotation measure

It has now been established that the Faraday rotation angle, $\Delta\psi$, is half of the phase difference between the right and left circularly polarised modes:

$$\begin{aligned}\Delta\psi &= \frac{1}{2}\Delta\phi \\ &= \frac{1}{2}\int_{s_0}^{s_1}\frac{\omega_\mathrm{e}^2|\Omega_\mathrm{e}|}{\omega^3}k\,\mathrm{d}r \\ &= \left(\frac{e^3}{2\pi m_\mathrm{e}^2c^4}\int_{s_0}^{s_1} n_\mathrm{e}B_\parallel\,\mathrm{d}r\right)\lambda^2.\end{aligned} \tag{4.13}$$

The bracketed term is defined as the rotation measure (RM):

$$\mathrm{RM} \equiv \frac{e^3}{2\pi m_\mathrm{e}^2c^4}\int_{s_0}^{s_1} n_\mathrm{e}B_\parallel\,\mathrm{d}r, \tag{4.14}$$

where $B_\parallel$ is the magnetic field component in the line of sight, making Equation (4.13)

$$\Delta\psi = \mathrm{RM}\,\lambda^2. \tag{4.15}$$

4.1.3 Polarised Foregrounds in Stokes Parameters

The signal measured by a radio telescope is a convolution of the emission from the sky and the instrument's own response. When considering polarised foregrounds, a complete mathematical framework is required to describe both the incoming radiation and the antenna's specific sensitivity to its polarisation state. This is achieved using the Stokes parameters to describe the sky and the Jones matrix formalism for the instrument.

The Stokes parameters

The full polarisation state of an electromagnetic wave can be described by a set of four intensity parameters called the Stokes parameters. These are the total

intensity, I, two parameters describing the state of linear polarisation, Q and U, and one for circular polarisation, V. For polarised synchrotron emission with an intrinsic polarisation angle ψ, the emission coefficients for the Stokes Q and U parameters are related to the total intensity emission coefficient, j_{b}^{I}, by:

$$j_{\mathrm{b}}^{Q} = j_{\mathrm{b}}^{I} \cos 2\psi, \tag{4.16}$$

$$j_{\mathrm{b}}^{U} = j_{\mathrm{b}}^{I} \sin 2\psi. \tag{4.17}$$

The brightness temperature, T_{b}, is then the integral of these emission coefficients along the line of sight, r:

$$T_{\mathrm{b}}(\nu) = \int j_{\mathrm{b}}(\nu)\, \mathrm{d}r, \tag{4.18}$$

From the final measured Stokes parameters, the total linearly polarised intensity and the observed polarisation angle, ψ_{obs}, can be reconstructed:

$$I = \sqrt{(Q)^2 + (U)^2}, \tag{4.19}$$

$$\psi_{\mathrm{obs}} = \frac{1}{2} \arctan\left(\frac{U}{Q}\right). \tag{4.20}$$

Instrumental response and polarisation leakage

The Stokes parameters have a direct physical basis in the electric field of the incoming wave. For an electric field described by a pair of perpendicular linear components (E_x, E_y), the total intensity, Stokes I, and the linear polarisation component, Stokes Q, are defined by the mean power in each component:

$$I \equiv \langle E_x^2 \rangle + \langle E_y^2 \rangle, \;\; Q \equiv \langle E_x^2 \rangle - \langle E_y^2 \rangle. \tag{4.21}$$

This allows the power in each electric field component to be expressed in terms of the Stokes parameters:

$$\langle E_x^2 \rangle = \frac{I+Q}{2}, \;\; \langle E_y^2 \rangle = \frac{I-Q}{2}. \tag{4.22}$$

As it shows that for an instrument sensitive only to a single electric field component (e.g., E_x) will measure a signal proportional to the sum (e.g., $I+Q$). This mixing of the polarised component, Q, into the total intensity measurement, I, is known as polarisation leakage.

To formalise this, the response of an antenna to an arbitrarily polarised wave is described by its 2×2 Jones matrix, $\mathbb{J}$. For a single-receptor instrument sensitive to only one linear polarisation, such as the dipole antenna modelled in this work, the Jones matrix takes a simplified form:

$$\mathbb{J} = J \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{4.23}$$

The measurement equation

The full relationship between the intrinsic sky brightness, described by the Stokes vector $\mathbf{s} = (I, Q, U, V)^{\mathrm{T}}$, and the signal measured by the antenna's receptors, $\mathbf{e} = (\mathcal{E}_{xx}, \mathcal{E}_{yy})^{\mathrm{T}}$, is given by the general measurement equation:

$$\mathbf{e}(\hat{\mathbf{r}}, \nu) = [\mathbb{J}(\hat{\mathbf{r}}, \nu) \otimes \mathbb{J}^*(\hat{\mathbf{r}}, \nu)]\mathbb{S}\,\mathbf{s}(\hat{\mathbf{r}}, \nu), \tag{4.24}$$

where $\mathbb{S}$ is a matrix that transforms the Stokes parameters into a frame the Jones matrix can act on:

$$\mathbb{S} = \frac{1}{2}\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \end{pmatrix}. \tag{4.25}$$

By substituting the simplified Jones matrix for a single-receptor instrument (Equation 4.23) into this general formalism, the resulting signal measured by the single receptor, $\mathcal{E}_{xx}$, is confirmed to be proportional to a linear combination of I and Q:

$$\mathcal{E}_{xx}(\hat{\mathbf{r}}, \nu) = \frac{1}{2} J_x^2(\hat{\mathbf{r}}, \nu)[I(\hat{\mathbf{r}}, \nu, t) + Q(\hat{\mathbf{r}}, \nu, t)]. \tag{4.26}$$

This demonstrates the phenomenon of polarisation leakage, where a component of the polarised sky (Q in this case) contaminates the measurement of the total intensity signal (I). After integration over the antenna beam, the measured brightness temperature, T_{xx}, is therefore a sum of the true foreground temperature (T_{f}, from I), the contaminating temperature from polarisation leakage (T_Q, from Q), and the 21-cm signal (T_{21}):

$$\begin{aligned} T_{xx}(\hat{\mathbf{r}}_{\mathbf{0}}, \nu, t) &\equiv \frac{\int_\Omega \mathcal{E}_{xx}(\hat{\mathbf{r}}', \nu)\,\mathrm{d}\hat{\mathbf{r}}'}{\int_\Omega A_x(\hat{\mathbf{r}}', \nu)\,\mathrm{d}\hat{\mathbf{r}}'} \\ &= T_{\mathrm{f}}(\hat{\mathbf{r}}_{\mathbf{0}}, \nu, t) + T_Q(\hat{\mathbf{r}}_{\mathbf{0}}, \nu, t) + T_{21}(\nu), \end{aligned} \tag{4.27}$$

where $A_x = J_x^2/2$ represents the power pattern (normalised shape of the directivity) of the antenna.

4.1.4 Rotation Measure Synthesis Technique

The Faraday rotation measure (RM) synthesis technique is a powerful tool for analysing polarised emission when multiple sources or rotating regions exist along a single line of sight. It employs the Fourier relationship between the observed polarisation properties and a function describing the intrinsic polarisation, the Faraday dispersion function, to separate the polarised sources along a line-of-sight.

For the simple case of a single polarised source behind a single rotating medium, the observed polarisation angle has a linear relationship with the square of the wavelength:

$$\chi(\lambda^2) = \chi_0 + \psi \cdot \lambda^2, \tag{4.28}$$

where χ_0 is the intrinsic polarisation angle, and $\chi(\lambda^2)$ is the observed polarisation angle at wavelength λ. Faraday depth is represented by ψ; it is determined by the

physical properties of the intervening medium integrated along the line of sight towards the source (Burn, 1966):

$$\psi \propto \int_0^{r_{\rm source}} \mathrm{d}r n_e B_{\parallel}, \tag{4.29}$$

where n_e is the free electron density, $B_{\parallel}$ is the magnetic field component along the line of sight, and $\mathrm{d}r$ represents the infinitesimal path length. For this simple case, the Faraday depth is equivalent to the rotation measure, which can be defined as the gradient of the polarisation angle with respect to λ^2:

$$\mathrm{RM} = \frac{\mathrm{d}\chi(\lambda^2)}{\mathrm{d}\lambda^2}. \tag{4.30}$$

This definition describes how the rotation measure is calculated, while the previous Equation (4.14) describes the underlying physical source of the rotation.

However, the interstellar medium is complex; multiple sources with varying Faraday depths may exist along the same line of sight. To deconstruct this complexity, the Faraday dispersion function, $\tilde{P}(\psi)$, is introduced. This function represents the complex polarised surface brightness per unit of Faraday depth. It is related to the observable complex surface brightness, $\tilde{P}(\psi)$, through a Fourier transform pair (Burn, 1966):

$$P(\lambda^2) = \int_{-\infty}^{+\infty} \tilde{P}(\psi) e^{2i\psi\lambda^2} \,\mathrm{d}\psi, \tag{4.31}$$

and

$$\tilde{P}(\psi) = \int_{-\infty}^{+\infty} P(\lambda^2) e^{-2i\psi\lambda^2} \,\mathrm{d}\lambda^2. \tag{4.32}$$

Here, $P(\lambda^2)$ is physically meaningful only for $\lambda^2 \geq 0$. In practice, any real observation can only sample a finite range of λ^2. This observational limitation is described by a window function, $W(\lambda^2)$, introduced by Brentjens and de Bruyn (2005). The actually observed polarised emission, $P_{\rm obs}$, is therefore the true emission multiplied by this window function:

$$P_{\rm obs}(\lambda^2) = W(\lambda^2) P_(\lambda^2) = W(\lambda^2) \int_{-\infty}^{+\infty} F(\psi) e^{2i\psi\lambda^2} \,\mathrm{d}\psi. \tag{4.33}$$

Due to the convolution theorem, this multiplication in λ^2-space corresponds to a convolution in the conjugate Faraday depth space. The reconstructed Faraday dispersion function, $\tilde{P}_{\rm obs}(\psi)$, is the true function convolved with the Fourier transform of the window function:

$$\tilde{P}_{\rm obs}(\psi) = K \int_{-\infty}^{+\infty} P_{\rm obs}(\lambda^2) e^{-2i\psi(\lambda^2-\lambda_0^2)} \,\mathrm{d}\lambda^2 = \tilde{P}(\psi) * R(\psi). \tag{4.34}$$

The function $R(\psi)$ is known as the rotation measure spread function (RMSF). It acts like an instrumental point-spread function, smearing out the true Faraday

depth structure. The RMSF is defined as the Fourier transform of the window function:

$$R(\psi) \equiv K \int_{-\infty}^{+\infty} W(\lambda^2) e^{-2i\psi(\lambda^2-\lambda_0^2)} \, \mathrm{d}\lambda^2, \tag{4.35}$$

where K is a normalisation constant. In practice, with data from N discrete frequency channels, these integrals are replaced by sums:

$$\begin{aligned} \tilde{P}_{\mathrm{obs}}(\psi) &\approx K \sum_{c=1}^{N} P_{\mathrm{obs},c} e^{-2i\psi(\lambda_c^2-\lambda_0^2)}, \\ R(\psi) &\approx K \sum_{c=1}^{N} W_c e^{-2i\psi(\lambda_c^2-\lambda_0^2)}, \end{aligned} \tag{4.36}$$

where the index c represents the individual frequency channel. The output of the rotation measure synthesis technique is therefore an observed Faraday depth cube, which represents the true sky convolved with the RMSF.

4.2 Simulating Observation with Linearly Polarised Foregrounds

Observation made by an antenna is unavoidably polarised. The observed brightness temperature contains the intrinsic intensity as well as polarisation leakage, and its direction depends on the type and the pointing of the antenna. The leakage due to polarised foregrounds can be described by Stokes parameters, which are directly measurable quantities.

The total brightness temperature observed by an antenna would be the linear combination of the intrinsic foreground intensity (Stokes I) and the leakage. The leakage is the combination of Stokes Q, U, and V, and the composition depends on the type of the radiometer and the alignment between the polarised emission and the radiometer (Appendix 4.1.3).

There are currently no all-sky observation on Stokes Q, U, and V maps in the relevant frequencies, and the emission leakage due to polarisation is not well understood. The aim of this work is to be as general as possible, and simplify the simulation to have tunable parameters under control to test the pipeline. These models are used to generate a variety of Stokes parameter maps to simulate different polarisation leakage scenarios.

This section details the three polarised foreground models used in this work: first, a model based on data given by the MWA 2400 degree^2 survey; second, an all-sky Stokes map created from a catalogue of point source rotation measures; and third, a 'toy model' with randomly assigned Faraday depth values in different sky regions.

4.2.1 Simulating global 21-cm signal Observation

Based on the established model introduced in Section 2.4.1, the simulated sky temperature, $T_{\rm sky}$, consists of two components: the linearly polarised foregrounds, $T'_{\rm f}$ ($T_{\rm f} + T_Q$ where $T_{\rm f}$ is the total foreground intensity and, T_Q, is the linearly polarised component as defined in Equation 4.27), and the sky-averaged 21-cm signal, T_{21}:

$$T_{\rm sky}(\Omega, \nu) = T'_{\rm f}(\Omega, \nu) + T_{21}(\Omega, \nu). \tag{4.37}$$

The sky model is generated using Equation (2.13). The simulated Stokes Q map T_Q, which will be detailed in Section 4.2.2, is added to the plain foreground (Equation 4.27) to produce a foreground with polarisation leakage, $T'_{\rm f}$:

$$T'_{\rm f}(\Omega, \nu) = T_{\rm f}(\Omega, \nu) + T_Q(\Omega, \nu). \tag{4.38}$$

Including the injected sky-averaged 21-cm signal, approximated as a Gaussian in this work, the simulated sky temperature Equation (4.37) is rewritten as:

$$T_{\rm sky}(\Omega, \nu) = T_{\rm f}(\Omega, \nu) + T_Q(\Omega, \nu) + T_{21}(\Omega, \nu). \tag{4.39}$$

The contaminated sky temperature is convolved with the beam, B, and in this work, the REACH dipole beam specifically. An uncorrelated Gaussian noise term, $\hat{\sigma}$, is then added to account for random antenna temperature noise. The complete equation for the simulated antenna temperature data in a single time bin is then:

$$T_{\rm data}(\nu) = \frac{1}{4\pi} \int_0^{4\pi} B(\Omega, \nu) T_{\rm sky}(\Omega, \nu)\, {\rm d}\Omega + \hat{\sigma}. \tag{4.40}$$

The resulting $T_{\rm data}$ represents the simulated dipole antenna temperature, which includes polarisation leakage and the injected signal. Simulated ground losses are not included in the data. In this work, each dataset consists of 24 hours of data collected at 5-minute intervals over three nights (8 hours per night).

4.2.2 Simulating Stokes Parameter Sky Map

In this section, the three adopted models for all-sky Stokes Q map generation are detailed. The adopted models are motivated for introducing complexities that are under control: they are used to create extreme conditions by tuning quantified parameters to test the robustness of the adopted data analysis pipeline. It is expected that the effect from a realistic polarised sky will be milder than in many of the cases presented here.

To generate foregrounds with leakage due to linear polarisation, a simulated all-sky Stokes Q map is added to the Stokes I (total intensity) map. All simulations are performed over a 50 MHz to 130 MHz frequency range with 1 MHz resolution, matching the operational band of the REACH dipole antenna (Cumner et al., 2022).

Stokes Q map by rotation measure synthesis technique

The first model tested is a simulation of polarised Galactic diffuse emission from Spinelli et al. (2018), which excludes discrete point sources. This simulation is based on the Faraday rotation measure synthesis technique that generates all-sky Stokes Q maps at sky direction $\hat{n}$ and frequency ν. This model uses the MWA 2400 degree^2 survey at 189 MHz as its foundation. Because the original interferometric data was only sensitive to degree-scale structures, the simulation extrapolates this information to the larger, tens-of-degree scales relevant for global 21-cm experiments.

The Faraday rotation measure synthesis formalism (Equation 4.31) is used to reconstruct the simulated polarised sky at the required frequencies from polarised intensity data as a function of Faraday depth, ψ. In this work, the model excludes contributions from the higher Faraday depths, $\psi > 5$ rad m^{-2} (see the top panel in Figure 4.1). This choice is motivated by two factors. First, it is consistent with observations showing that diffuse polarised emission at low frequencies is primarily confined to low Faraday depths (Haverkorn et al., 2004; Bernardi et al., 2009; Lenc et al., 2016). Second, this creates a more challenging test case, as the resulting contamination has smoother spectral structures that are more difficult to distinguish from the global 21-cm signal.

Rotation measure map of point sources

The second model generates an all-sky Stokes Q map using a catalogue of observed point source rotation measures. Presented by Van Eck et al. (2023), RMTable2023 is a consolidated catalogue of rotation measures from point sources collected from 42 published catalogues. This catalogue is used to simulate another all-sky Stokes Q map. The catalogue itself consists of sparse data from discrete radio sources, so the data are first interpolated using nearest-neighbour interpolation to create a continuous, all-sky map of rotation measure. The middle panel in Figure 4.1 shows the interpolated all-sky map. The Stokes Q value in each pixel is calculated using the relation:

$$Q = Ip\cos(2\psi\lambda^2), \tag{4.41}$$

where I is the total intensity given by the scaled 2008 Global Sky Model (de Oliveira-Costa et al., 2008), and p is the linear polarisation fraction. The formalism of the equation is motivated by the λ^2 dependence in the Faraday rotation angle of synchrotron emission (Equation 4.28). The total intensity of an emission includes both polarised and unpolarised components; p, the linear polarisation fraction, is the ratio of these components, and is constrained by $0 \leq p \leq 1$.

Maps with different numbers of Faraday depth region

The third approach is a toy model designed to systematically investigate how the spatial complexity of the Faraday depth distribution affects signal extraction. In this model, the sky is divided into a set number of regions, and each is assigned a random Faraday depth value. The sky is divided in two different ways: the first one is a simple division into equally sized regions using the HEALPix framework (Górski et al., 2005): 12 ($N_{\text{side}} = 1$), 48 ($N_{\text{side}} = 2$), and 196 ($N_{\text{side}} = 4$) regions respectively. The selected numbers of sides are the lowest possible settings for HEALPix gridding. The second one is a physically motivated division based on the sky's spectral index structure based on a foreground modelling method developed for the REACH data analysis pipeline. The size of each region varies depending on the local spectral index distribution (Anstey et al., 2021). The bottom panels in Figure 4.1 show one example for each case. As with the previous model, the value of the Stokes Q in each pixel is calculated using Equation (4.41).

To create a challenging scenario, the randomly assigned Faraday depths are deliberately restricted to low values (< 0.5 rad m^{-2}). This ensures that the resulting polarisation leakage has a smooth, slow oscillation with frequency, making it spectrally similar to the global 21-cm signal and thus more difficult for the pipeline to separate.

4.3 Results

In this chapter, a recovered signal with signal RMSE (Equation 2.21) lower than 30% of the injected signal strength is considered a good fit with low signal RMSE. The results are divided into three sections based on the three different models introduced in Section 4.2.

4.3.1 Rotation Measure Synthesis Technique

Spinelli et al. (2019) showed that the oscillatory behaviour of the Stokes Q parameter as a function of frequency reveals notable differences between cases in which all Faraday depths are taken into account ('all ϕ') and those restricted to low Faraday depth values ($\psi \leq 5$ rad m^{-2}, 'low ϕ'). In the low-frequency band (≤ 100 MHz), oscillatory structures arising from polarisation leakage are more complex and exhibit greater contamination overall in the 'all ϕ' cases than in the 'low ϕ' cases. Conversely, in the 'low ϕ' cases, the oscillations become smoother in the high-frequency band (> 100 MHz) as well as slower and larger in magnitude.

To investigate how the results depend on the signal's position relative to the foreground's oscillatory structure, tests were performed with injected signals centred at both 90 MHz and 120 MHz with an amplitude A of 0.157 K and $\sigma = 15$ MHz. In the cases where all Faraday depth values were included, the

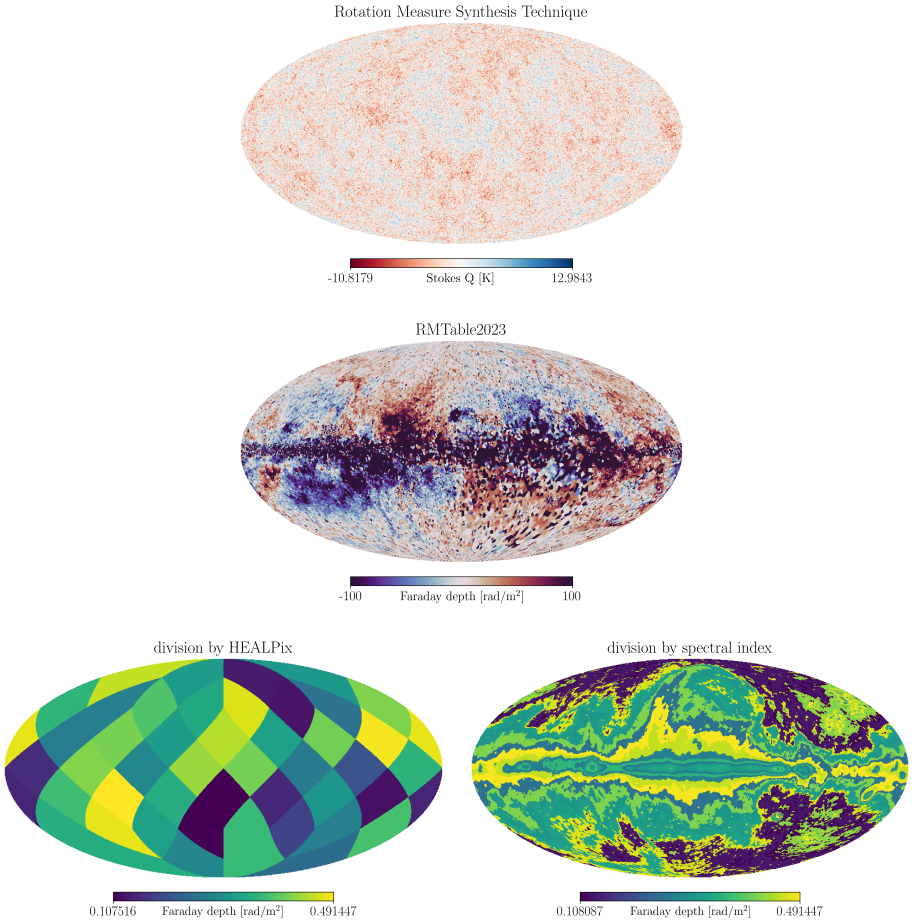


Figure 4.1: The three different models used to simulate all-sky polarisation leakage, shown in Galactic coordinates. **Top:** A Stokes Q map generated from a simulation of diffuse Galactic emission (Spinelli et al., 2018) that is limited to low Faraday depths ($\psi < 5$ rad m^{-2}) as observations have shown that polarised emission in the low-frequency band has little emission at high Faraday depths (Haverkorn et al., 2004; Bernardi et al., 2009; Lenc et al., 2016). It yields a relatively slower oscillation with respect to frequency compared to cases where all Faraday depths are taken into account. **Middle:** An all-sky Faraday depth (ψ) map created by interpolating the RMTable2023 catalogue of discrete point sources (Van Eck et al., 2023). The catalogue itself consists of data from discrete polarised radio sources and does not cover the whole sky, so the Faraday depth of the areas without data is interpolated using the nearest-neighbour interpolation. **Bottom:** Two toy models showing sky divisions with randomly assigned low Faraday depths (< 0.5 rad m^{-2}). **bottom-left:** A division into 48 regions ($N_{side} = 2$) using the HEALPix scheme. **bottom-right:** A division into 25 regions based on the sky's spectral index structure.

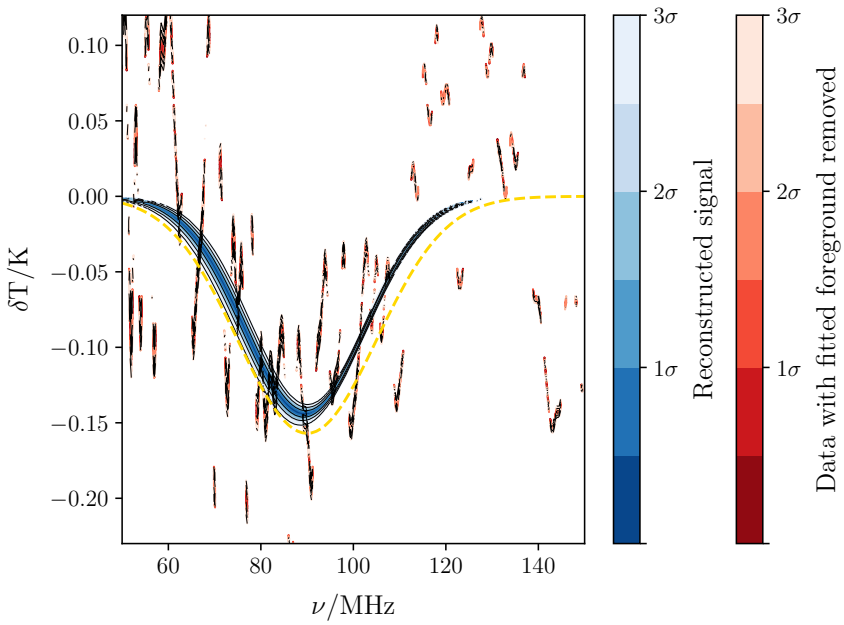


Figure 4.2: The reconstructed signal (blue contours) and foreground residual (red contours) plots of an example case by the rotation measure synthesis technique. The yellow dashed line represents the injected signal with an amplitude of 0.157 K centred at 90 MHz. This example considers only low Faraday depth values ($\psi \leq 5$ rad m^{-2}). Visually, the reconstructed signal contours are already very close to the injected signal. The reconstructed signal RMSE is 0.0138 K, which is much lower than 30% of the injected signal strength.

position of the signal had no significant effect on the reconstructed signal itself; however, the magnitude of the foreground residuals did differ significantly between the two scenarios. This level of contamination would already pose a challenge for most other standard pipelines; the analysis in Spinelli et al. (2018) required a reduction in the contamination amplitude to achieve a successful signal recovery.

In the 'low ϕ' cases, the results are more sensitive to the signal's centre frequency. For a signal at 90 MHz, the signal recovery is robust, as the contamination in the low-frequency part of the band is minimal. However, when the signal was centred at 120 MHz, a significant increase in the foreground residuals was observed. This is due to two effects. First, the 'low ϕ' contamination at these higher frequencies oscillates at a rate spectrally similar to the global 21-cm signal, making the two more difficult to separate. Second, because the intrinsic foreground brightness is much lower at high frequencies, the relative impact of the contamination is more prominent.

The time-separated pipeline is able to recover the signal with low signal RMSE (lower than 30% of the injected signal strength) and high log Bayes ratio ($\Delta \log \mathcal{Z}$) in all the tested cases. Figure 4.2 shows one typical example of a 'low ϕ' case.

The major distinction between the 'all ϕ' and 'low ϕ' scenarios lies in the spectral character of the resulting contamination. In the 'all ϕ' case, the superposition of many Faraday depth components produces a rapidly oscillating spectral structure that is spectrally distinct from the smooth, broad absorption feature of the global 21-cm signal. In contrast, the 'low ϕ' cases, by construction, generate spectrally smoother contamination that evolves more slowly with frequency.

While the time-separated pipeline recovers the signal in both scenarios, the 'low ϕ' case is more insidious. Its smoother spectral form presents a greater challenge for foreground separation, as it more closely resembles the signal shape. A standard analysis pipeline, as noted in Spinelli et al. (2018), would struggle with this level of contamination. This underscores that the primary risk from polarised foregrounds is not necessarily leakage with high amplitude, but leakage that is spectrally simple. Such contamination, arising from large-scale regions of low Faraday depth, is a potential source of a false positive detection.

4.3.2 Rotation Measure Catalogue of Point Sources

This model, based on a catalogue of observed point sources, serves to quantify the tolerance of the analysis to contamination from a realistic, spatially complex population of discrete sources. The all-sky Stokes Q map used in this section is simulated based on the RMTable2023 catalogue given in Van Eck et al. (2023) via Equation (4.41). The oscillation of the Stokes Q parameter as a function of frequency, shown in Figure 4.3, appears noisy. The noisy oscillation can be explained by the distribution of many different Faraday depth values over the entire sky; the superposition of $\sin(a\lambda^2)$ waves oscillating at different rates will result in seemingly noisy structures. There are also a significant amount of regions with high Faraday depth values (middle panel in Figure 4.1), which by Equation (4.41) would intrinsically result in fast oscillation.

The limited bandwidth of the dipole beam restricts the number of different centre frequencies at which the injected signal can be tested. Signals centred in the regions near the frequency band edges are difficult for the pipeline to constrain, often leading to an overestimation of their depth.

Figure 4.4 shows the results of tests using uniform linear polarisation fractions (p) of 1/10, 1/30, and 1/70. These values were chosen to probe the pipeline's breaking point, as the observed polarisation fraction for such sources is typically low (the polarisation fraction for the radio sources at 189 MHz have an upper limit of $\sim 2\%$ (Bernardi et al., 2013)). The injected signal has an amplitude of 0.157 K, a width of 15 MHz and is centred at 90 MHz. For the extreme case of $p = 1/10$, the REACH pipeline fails to recover the injected signal. Whereas for the lower fractions, the signal is successfully recovered, albeit with a slightly diminished amplitude.

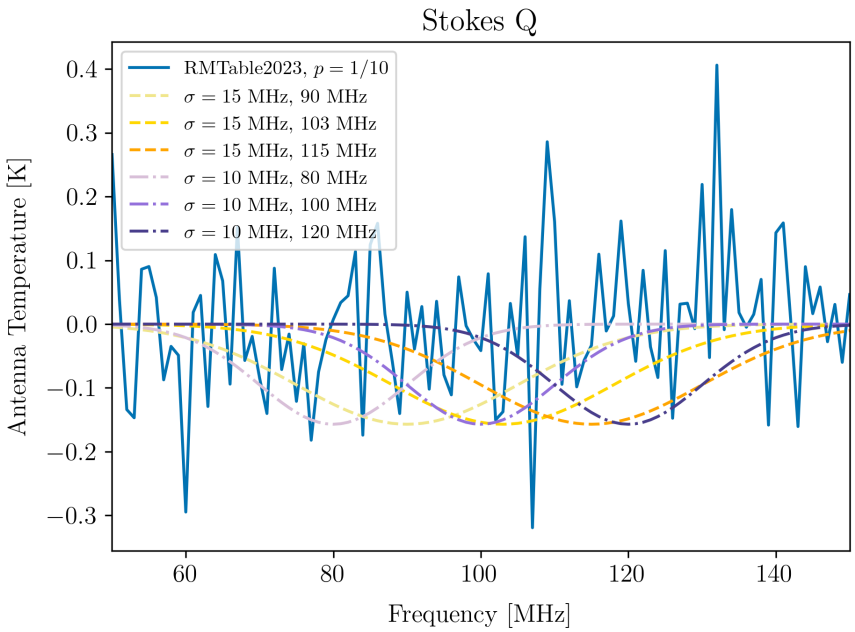


Figure 4.3: The scaled Stokes Q as a function of frequency. The model is simulated (Equation 4.41) using the RMTable2023 catalogue (Van Eck et al., 2023) with a uniform linear polarisation fraction of $p = 1/10$. For comparison, two smooth, injected global 21-cm signals are also shown, with widths of 15 MHz (yellow dashed line) and 10 MHz (purple dashed line). The plot illustrates the spectrally complex, noisy character of the contamination compared to the simple shape of the global 21-cm signal.

To test for any frequency-dependent susceptibility, the analysis is repeated for signals centred at 80 MHz, 100 MHz, and 120 MHz. As shown in Figure 4.5, the signal recovery improves significantly with decreasing polarisation fraction across all frequencies tested. Signal RMSE of less than 10% of the injected signal absorption strength can be achieved for linear polarisation fractions of 1/30 and 1/10. The results do not suggest that any particular frequency range is more susceptible to this type of polarisation contamination, as discussed further in Appendix 4.A.

This establishes a quantitative requirement for global 21-cm experiments. Given that the observed polarisation fraction of sources is typically low, the findings suggest that contamination from discrete sources is manageable provided the instrumental polarisation response is well-characterised and stable.

4.3.3 Number of Faraday Depth Regions

This toy model provides a controlled test of how the spatial complexity of the Faraday depth distribution impacts signal extraction. The first set of cases uses a sky division scheme based on spectral index similarity (as shown on the right-hand side of Figure 4.1),the same scheme used for foreground modelling in the

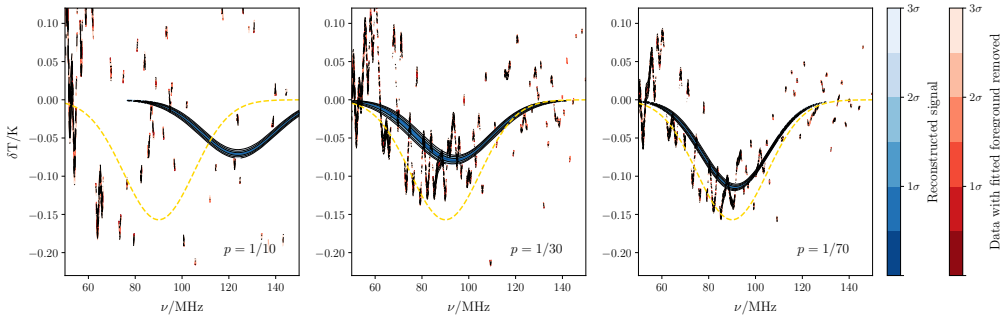


Figure 4.4: Reconstructed signal (blue contours) and foreground residual (red contours) for cases with different uniform linear polarisation fractions (p) of 1/10, 1/30, and 1/70 from left to right. The same polarisation fraction is applied to the whole sky at all frequencies in each case. The Faraday depth map (middle panel in Figure 4.1) used to generate the data is based on the RMTable2023 catalogue (Van Eck et al., 2023). The injected signal is marked by the yellow dashed line. The figure clearly demonstrates that signal reconstruction improves significantly as the polarisation fraction decreases. The model includes a term for random noise but it has an upper limit; if the noise is bigger than the upper limit, it would not be completely removed, which is manifested by the noisy residuals in the red contours.

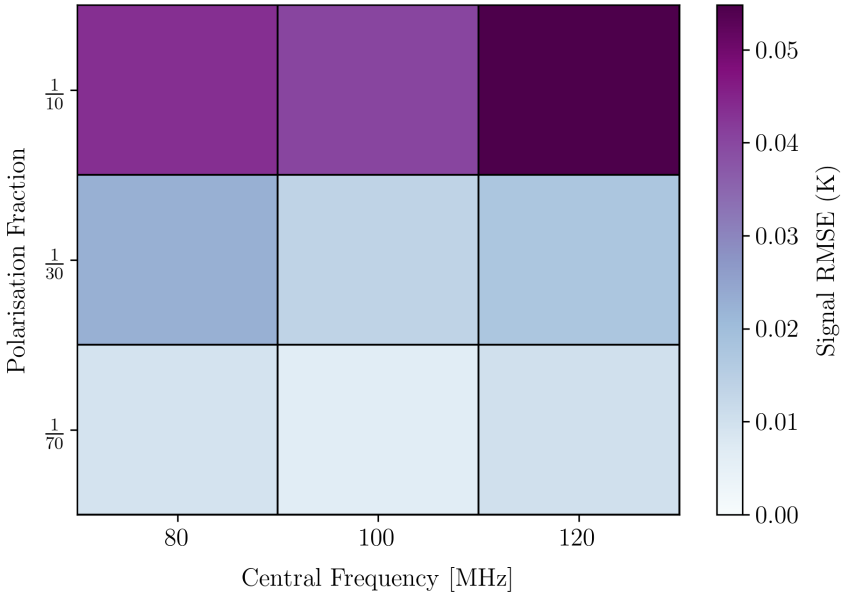


Figure 4.5: Signal RMSE as a function of linear polarisation fraction and injected signal centre frequency for the point source model. The injected signals have an amplitude of 0.157 K and a width of 10 MHz, and are centred at 80 MHz, 100 MHz, and 120 MHz respectively. The grid shows a strong trend of improving signal recovery (lower RMSE) with decreasing polarisation fraction, but no significant trend with the signal's centre frequency. Further details are discussed in Appendix 4.A.

REACH pipeline. In Figure 4.6, Stokes Q is shown as a function of frequency with a uniform linear polarisation fraction[a] of $p = 1/1000$ across the whole sky at all frequencies. This value is deliberately chosen to generate a contamination signal with a spectral structure comparable in magnitude to the expected global 21-cm signal, creating a challenging test case. Each curve corresponds to a case with a distinct number of Faraday depth regions obtained by spatially integrating over the entire Stokes Q parameter maps simulated at each frequency according to Equation (4.41). The oscillation of Stokes Q as a function of frequency is slower and smoother when the number of regions is low, with its overall magnitude being comparable to the injected global 21-cm signal.

Figure 4.7 shows the signal RMSE of cases with different numbers of regions and different polarisation fractions. Two clear trends are evident. First, as expected, signal recovery improves as the polarisation fraction decreases, as the magnitude of the contamination has a linear dependency on the polarisation fraction (Equation 4.41). This trend of improving recovery with more regions arises because a larger number of regions, each with a random Faraday depth, produces a sky-averaged contamination signal that is spectrally more complex and faster-oscillating. This analysis reveals that the spectral character of the contamination has a greater impact on signal recovery than its overall magnitude. For example, the case with 36 regions and $p = 1/1000$ yields a significantly lower RMSE than the case with only 2 regions and $p = 1/3000$, even though the total contamination amplitude is similar in both scenarios. Figure 4.8 shows visual examples of the reconstructed signals for several cases. The division by HEALPix gridding scheme produces similar results to the spectral-index based division of more than 12 regions, where the signal RMSE has saturated, as the lowest allowed grid number (12 regions) is already high.

This result shows that a spatially simple polarised foreground, dominated by a few low-Faraday-depth regions, presents a greater challenge to a global 21-cm experiment. A more complex sky, with a more varied Faraday depth structure, is advantageous, as the beam integration averages down the polarised signal into a more spectrally complex and lower-amplitude contaminant. This reinforces that the spectral character of the contamination, dictated by the large-scale Faraday structure of the Galaxy, is more critical than its intrinsic, small-scale amplitude.

[a]The fractions in this model are chosen to be much lower than for the point source models discussed previously, reflecting line-of-sight depolarisation (Burn, 1966). While individual sources can be polarised at the percent level, integrating over many regions with randomly oriented magnetic fields reduces the net polarisation to a much lower value. The specific values are selected to generate a challenging test case where the resulting polarisation leakage is comparable in magnitude to the expected 21-cm signal, thereby rigorously testing the pipeline's capabilities.

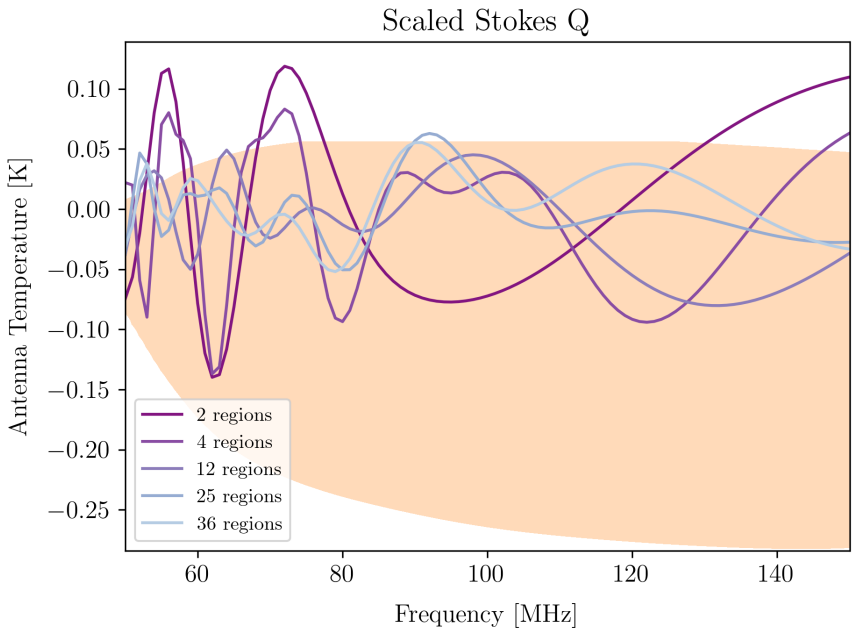


Figure 4.6: Scaled Stokes Q as a function of frequency for different numbers of Faraday depth regions, with a uniform linear polarisation fraction of $p = 1/1000$ across all frequencies for all cases. The regions are divided according to the spectral indices using the REACH data analysis pipeline (Anstey et al., 2021). The light-orange shading represents the prior range for a standard global 21-cm signal within 3 sigmas, and the solid lines show the Stokes Q oscillation for cases divided into different numbers of regions. The Stokes Q parameter map is simulated according to Equation (4.41). The oscillation of Stokes Q is slower and smoother when the number of regions is low, with its magnitude being comparable to the injected global 21-cm signal.

4.4 Conclusions

Contamination from a polarised sky poses a potential challenge to global 21-cm experiments. In this chapter, three different approaches are adopted to simulate foregrounds with linear polarisation leakage observed by a dipole beam. Simulated effective antenna temperature data, including an injected global 21-cm signal, are analysed using REACH's time-separated data analysis pipeline, which jointly fits data across all time bins.

The first approach of polarisation modelling employs rotation measure synthesis technique to generate an all-sky Stokes Q sky map for each frequency in the observed range (50 to 150 MHz with 1 MHz resolution) along the line of sight. The model is based on the extrapolation from the MWA 2400 degree^2 survey at 189 MHz and simulates the polarised Galactic diffuse emission. The resulting contamination structure agrees with those reported by Spinelli et al. (2018) and Spinelli et al. (2019), that 'low ϕ' contamination oscillates at a rate that makes it resemble the global 21-cm signal better, and therefore makes it more challenging

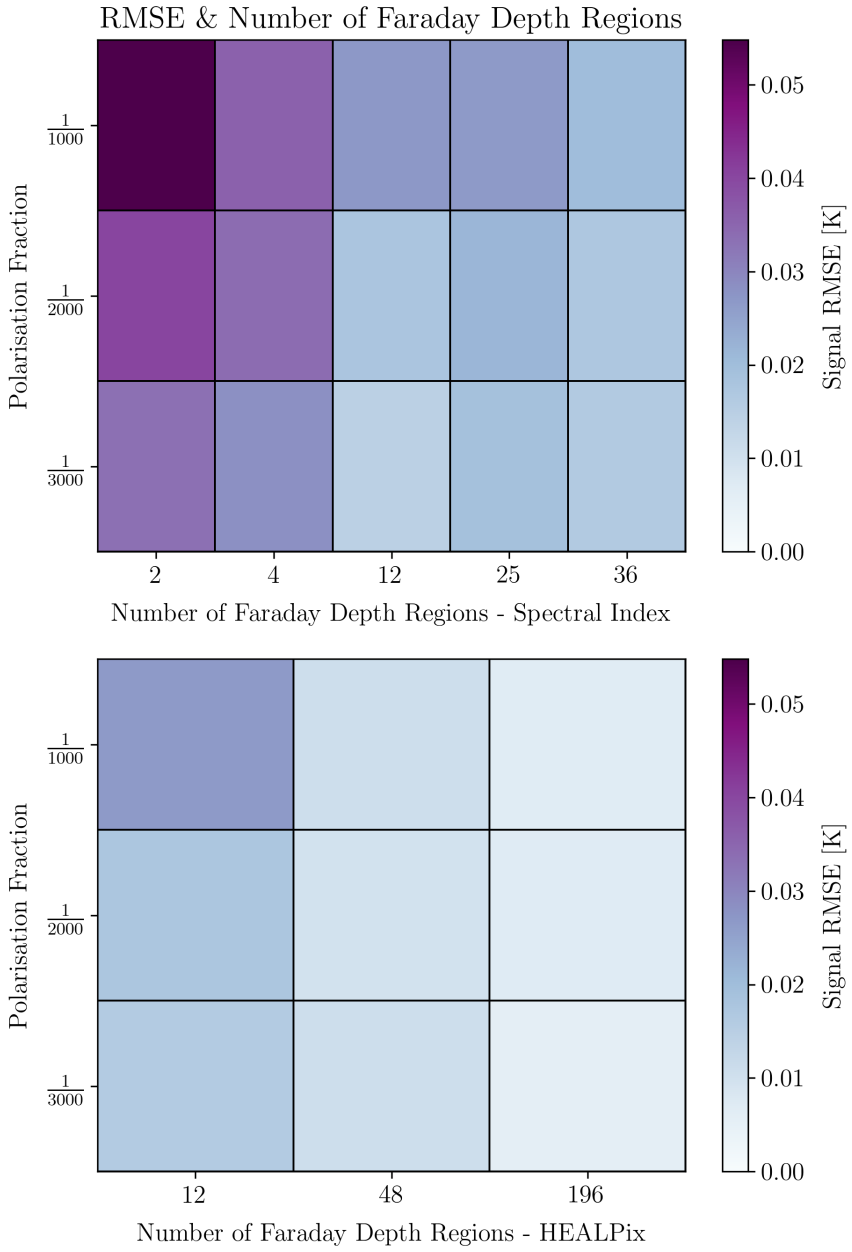


Figure 4.7: Signal RMSE as a function of polarisation fraction and the number of Faraday depth regions used in the toy model. The injected signal in all cases has an amplitude of 0.157 K, a width of 15 MHz, and is centred at 90 MHz. The **top panel** shows results for a sky division based on spectral index structure (Anstey et al., 2021), while the **bottom panel** uses a HEALPix gridding scheme. The RMSE decreases as the number of regions increases because the resulting spectrally complex (noisier) contamination becomes easier to distinguish from the smooth signal. Since the minimum HEALPix division tested is 12 regions, a number at which the signal RMSE has already saturated, its results are comparable to the spectral index scheme in this high-region-count regime.

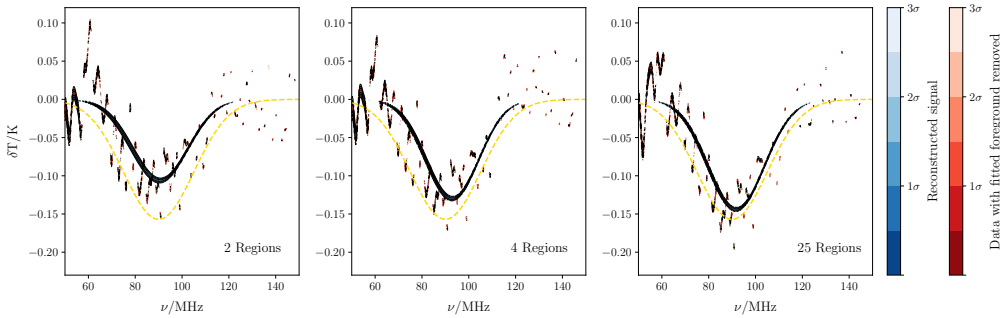


Figure 4.8: Reconstructed signal (blue contours) and foreground residual (red contours) for the toy model with 2, 4, and 25 Faraday depth regions (from left to right). In these cases, the sky is divided based on spectral index structure, and a uniform linear polarisation fraction of $p = 1/3000$ is applied. The Faraday depths are randomly assigned from 0 to 0.5 rad m^{-2}. The injected signal is marked by the yellow dashed line. The plots clearly demonstrate that signal recovery improves as the number of regions increases.

to separate. On the other hand, 'all ϕ' cases exhibit fast oscillatory structures that are easier to separate. Nevertheless, the time-separated REACH data analysis pipeline successfully distinguishes the signal in both 'low ϕ' and 'all ϕ' cases, achieving low signal RMSE and high log Bayes ratios. In Spinelli et al. (2018) and Spinelli et al. (2019), most other pipelines struggled under the same amount of contamination, requiring a reduction in the contamination amplitude to achieve results of similar accuracy.

The second approach uses the rotation measure data of discrete point sources from the RMTable2023 catalogue (Van Eck et al., 2023) to generate an interpolated all-sky Faraday depth map using nearest-neighbour interpolation. Three Stokes Q sky maps, each with a different linear polarisation fraction applied uniformly across the sky at all frequencies, are generated to model foreground maps with varying levels of polarisation contamination. The resulting Stokes Q maps have fast oscillatory structures. For injected signals centred between 80 MHz and 120 MHz, signal recovery by the time-separated REACH pipeline improves with decreasing polarisation fraction. A signal RMSE of less than 10% of the injected signal strength is achievable for linear polarisation fractions of $\leq 3\%$. Bernardi et al. (2013) found that polarisation fraction for the radio sources at 189 MHz in their catalogue has an upper limit of $\sim 2\%$. This suggests that the time-separated REACH data analysis pipeline is capable of separating the global 21-cm signal in the presence of polarisation contamination due to discrete radio sources.

The third approach uses a toy model that divides the sky into regions with randomly assigned low Faraday depth values to test the impact of complex Faraday depth distributions. Linear mixing of contamination oscillating at different rates across frequencies produces unsmooth, relatively noisy structures compared to the

foreground. The results indicate that increased variation and spatial complexity in Faraday depth values produce structures more distinct from the global 21-cm signal, making contamination easier to separate.

The time-separated REACH pipeline demonstrates robustness against these contaminants with injected global 21-cm signals of an amplitude of 0.157 K centred between 80 MHz and 120 MHz in the presence of unaccounted-for linear polarisation contamination in the majority of cases, particularly in scenarios that would challenge standard analysis methods. However, the results suggest a physical criterion for the validation of 21-cm detection: a candidate signal should be tested against the possibility that it arises from spectrally simple polarised contamination. This work provides a framework for understanding the nature of polarised contamination and quantifies the conditions under which it can be successfully mitigated.

4.A Difference between Two Orthogonally Pointing Radiometers

Figure 4.9 shows the signal RMSE for cases with injected signal centred at different frequencies using two radiometers whose directions of pointing are orthogonal to each other. The initial purpose is to determine whether certain frequency ranges are more susceptible to fast-oscillating contamination. Both $I-Q$ and $I+Q$ cases are analysed to observe comparable effects in both directions. $I-Q$ and $I+Q$ are equivalent to two radiometers orthogonal to each other. The REACH data analysis pipeline is known to underestimate, rather than overestimate, the depth of the global 21-cm signal when using the dipole beam in cases where foregrounds were inadequately modelled. This mitigates signal overestimation in cases where overall contamination enhances the absorption feature of the global 21-cm signal, as reflected in the results: the signal RMSE is relatively lower than its underestimated counterpart when it is overestimated that the relative degree of accuracy in signal recovery is almost the opposite between the $I-Q$ and $I+Q$ cases. This behaviour suggests that underestimation or overestimation is the primary driver of the relative signal RMSE difference, rather than any specific frequency range being particularly susceptible to polarisation contamination.

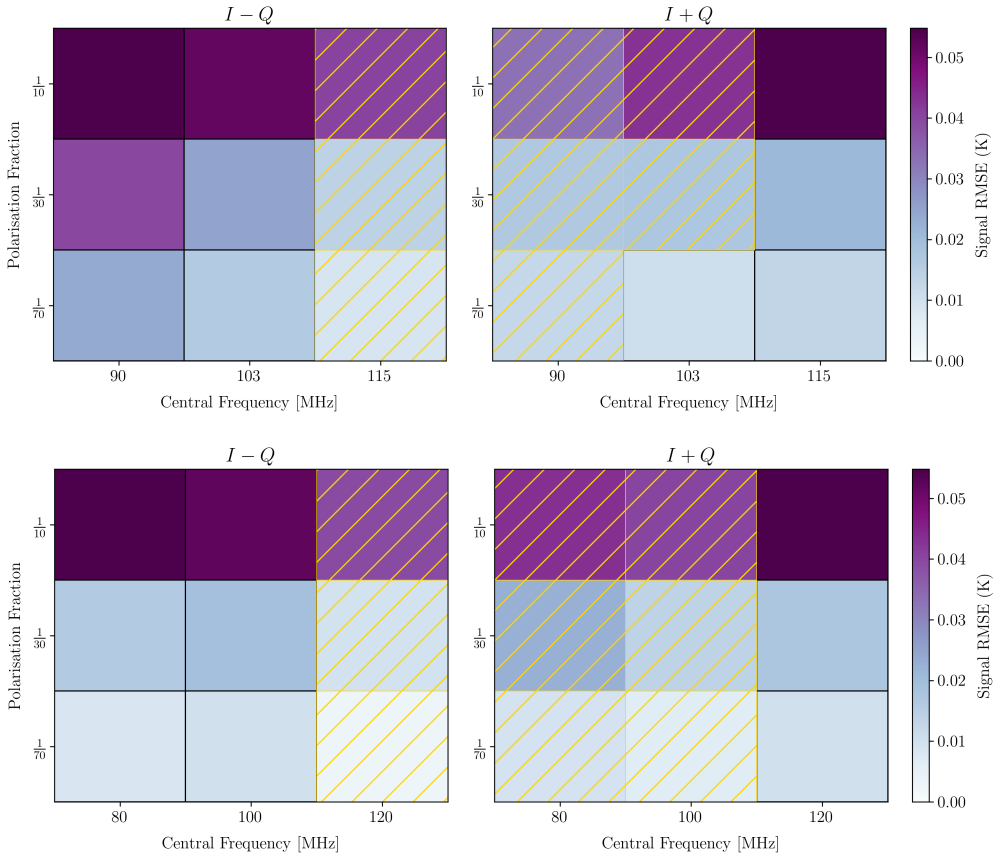


Figure 4.9: Signal RMSE as a function of linear polarisation fraction for injected signals with different parameters, simulated for two orthogonal radiometers ($I + Q$ on the left, $I - Q$ on the right). The **top row** uses a signal width of 15 MHz, centred at different frequencies: 90 MHz, 103 MHz, and 115 MHz while the **bottom row** uses a 10 MHz width, centred at different frequencies: 80 MHz, 100 MHz, and 120 MHz. Yellow hatches indicate cases where the signal amplitude was overestimated. As expected, signal recovery improves significantly as the polarisation fraction decreases. No significant difference in signal recovery is observed between injected signals centred at different frequencies or those with differing widths (see Appendix 4.A for a more detailed explanation).

Part III

Global Signal Model

Chapter 5

FlexKnot Signal Model

O time, thou must untangle this, not I!
It is too hard a knot for me to untie!

W. Shakespeare

The research presented in this chapter has been published in the Monthly Notices of the Royal Astronomical Society as Shen et al. (2024).

5.1 Introduction

As mentioned in Section 1.2.1, the exact evolution of the 21-cm signal is dependent on the timing, amplitude, and rate of defining astrophysical events in cosmic history. An example of the simulated signal by globalemu (Bevins et al., 2021) is shown in Figure 5.1: during the dark ages, gas is sufficiently dense for collisions to couple the spin temperature of the hydrogen gas to the kinetic temperature of the gas, $T_{\rm S} \to T_{\rm K}$. Due to the adiabatic cooling of gas, the kinetic temperature of the gas is lower than the radio background temperature, $T_{\rm K} < T_\gamma$. This difference results in an absorption against the radio background, making $T_{\rm S} < T_\gamma$.

The collisional transitions become negligible as gas density decreases due to continuing expansion, setting $T_{\rm S} \to T_\gamma$. During cosmic dawn, the first luminous objects begin to emit radiation in the Lyman-band. The spin temperature is then coupled to cold gas, $T_{\rm S} \sim T_{\rm K} < T_\gamma$ via Wouthuysen-Field effect (Wouthuysen, 1952; Field, 1958), resulting in another absorption.

Fluctuations as well as total intensity in the Lyman-α background no longer affect the 21-cm signal after Lyman-α coupling saturates. The growing population of X-ray sources start heating the adiabatically cooling gas, driving $T_{\rm S} \sim T_{\rm K} > T_\gamma$

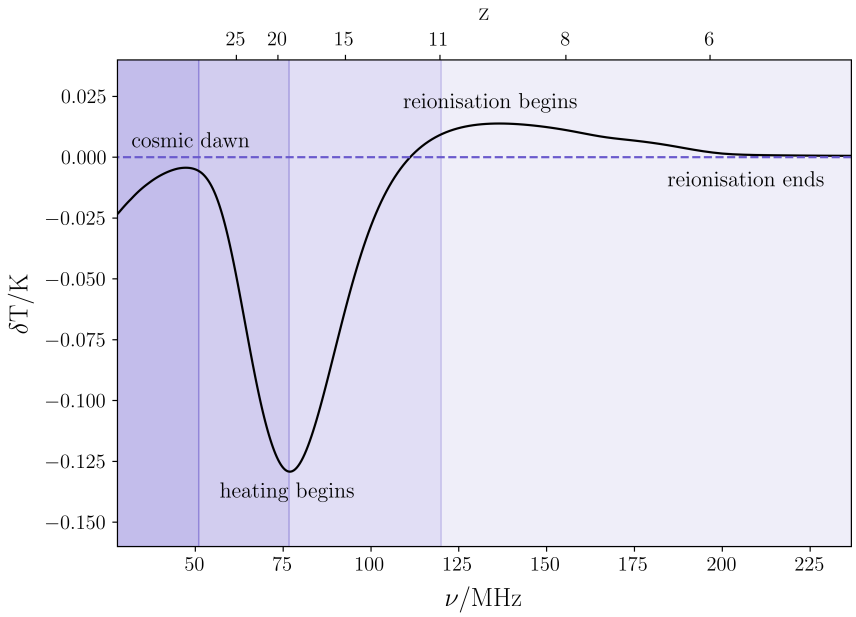


Figure 5.1: A simulated global 21-cm signal generated using GLOBALEMU (Bevins et al., 2021) with the following parameters: star formation efficiency $f_* = 0.02$, minimal virial circular velocity $V_c = 16.5$ km/s, X-ray efficiency $f_x = 1$, CMB optical depth $\tau = 0.06$, slope of the spectral energy density $\alpha = 1$, low energy cut-off of the X-ray spectral energy density $\nu_{\min} = 0.2$ keV, mean free path of the ionising photons $R_{\rm mfp} = 30$ Mpc. The major cosmic events are marked on the plot, namely, cosmic dawn (the onset of star formation), heating, and the beginning and the end of reionisation.

in general. In some models, this would turn absorption into emission. Ionising photons begin to turn neutral hydrogen (HI) to ionised hydrogen (HII), and then the signal eventually tends to zero (Shaver et al., 1999; Furlanetto, 2016).

Knowing the location and the shape of the global 21-cm signal absorption trough can enlighten us on the timing of the primordial star formation, the mass and star formation efficiency of the first star-forming halos, and the luminosity of the first X-ray sources (Furlanetto et al., 2006; McQuinn and O'Leary, 2012).

Existing methods of data analysis often assume a specific signal shape, such as a Gaussian, or the flattened Gaussian motivated by EDGES result (Bowman et al., 2018), which are unable to simultaneously describe both absorption and emission peaks. While computationally convenient, these approaches introduce a strong model-dependent bias: one can only find a signal that resembles the assumed template. On the other hand, physical models from emulators such as those of 21cmFAST (Mesinger et al., 2011), 21cmGEM (Cohen et al., 2020), and GLOBALEMU (Bevins et al., 2021) explore a wide parameter space but may not capture all possible signal morphologies, particularly if exotic physics or unexpected radio backgrounds are at play.

The observational state of the field, currently shaped by the conflicting results

of EDGES and the SARAS 3 experiment, illustrates the challenge of this model uncertainty. This motivates the need for a signal model that is sufficiently flexible to capture a wide range of astrophysically plausible shapes without being overly susceptible to fitting foreground residuals or instrumental systematics.

This chapter implements and tests the FlexKnot model, a signal model designed to be sufficiently flexible to describe a variety of shapes while remaining separable from foregrounds and systematics. This is realised by adopting a function parametrised by freely moving knots strung together by splines. FlexKnot allows the data to determine the signal's shape in a physics-agnostic manner, giving it the potential to recover features that may arise from unexpected physics.

The FlexKnot method was introduced by Vázquez et al. (2012) to reconstruct the primordial power spectrum from the Planck data, and then adopted by Millea and Bouchet (2018) for parameterising the reionisation history. Heimersheim et al. (2022) used it to parametrise the EoR history and the potential of hypothetical high-redshift Fast Radio Bursts to constrain cosmic reionisation.

Heimersheim et al. (2023) have also adopted the FlexKnot method as a global signal model. They used FlexKnot to separate the non-foreground component using the EDGES low-band data. There are several differences between this implementation and the approach of Heimersheim et al. (2023). In contrast to the work presented here, they do not distinguish between the global 21-cm signal and systematics. They modelled the foregrounds using a polynomial, while the foregrounds in this work are simulated with a physically motivated method. Moreover, instead of adopting Piecewise Cubic Hermite Interpolating Polynomial (PCHIP) to perform interpolation, this model starts by parameterising the second derivative of the global 21-cm signal via linearly interpolated splines, recovering the final signal via double integration.

The FlexKnot signal model is tested by implementing it in the REACH data analysis pipeline (Anstey et al., 2021) where the physically motivated foreground and the signal are jointly fitted using the Bayesian nested sampling algorithm PolyChord (Handley et al., 2015). The FlexKnot model is also tested on cases in which uncalibrated sinusoidal systematic structures are present. Unaccounted-for systematics can potentially arise in practice, due to e.g. ground plane artefact (Bradley et al., 2019) and calibration issues (Sims and Pober, 2020), so it is crucial to know how susceptible the signal model is to the presence of systematics.

5.2 Constructing the FlexKnot Function

The FlexKnot model in this work is based on Vázquez et al. (2012), who developed it to reconstruct the primordial power spectrum from Planck data. Due to its higher degree of flexibility, the model has been widely adopted in other studies (Bridges et al., 2009; Vázquez et al., 2012, 2013; Aslanyan et al., 2014; Abazajian

et al., 2014; Hee et al., 2016; Planck Collaboration et al., 2016; Hee et al., 2017; Olamaie et al., 2018; Millea and Bouchet, 2018; Planck Collaboration et al., 2020,; Heimersheim et al., 2022; Escamilla and Vazquez, 2023). In this work, it is formed by a number of knots interpolated by splines. It has the advantage of describing functions of all possible shapes, given a sufficient number of knots; a desirable feature given the uncertain shape of the global 21-cm signal. Normally, there is a deep absorption trough at $z \sim 10 - 30$ (corresponding to $\sim 45 - 130$ MHz) and an emission at $z \sim 6 - 10$ (corresponding to $\sim 130 - 200$ MHz); but this is model-dependent. Piecewise cubic spline interpolation is preferred for modelling the global 21-cm signal. In this work, the FlexKnot function is constructed by integrating twice over a function of freely moving knots interpolated by piecewise linear splines:

$$f_{\rm cubic}(\nu) = \int_{\nu_0}^{\nu} {\rm d}\nu' \int_{\nu_0'}^{\nu'} {\rm d}\nu'' f_{\rm linear}(\nu''), \tag{5.1}$$

which is equivalent to a function made up of piecewise cubic splines, and works around the difficulties in interpreting the posteriors and setting priors when directly employing cubic spline interpolation with nested sampling (Handley et al., 2019). The positions of the knots are determined by Bayesian nested sampling. Figure 5.2 shows an example of a signal recovered by the two different methods: the upper three panels show how the FlexKnot function is implemented as a signal model in three steps, and the last panel shows the direct cubic spline interpolation done via Piecewise Cubic Hermite Interpolating Polynomial (PCHIP). The number of knots required to recover a given signal shape, such as a Gaussian, differs between the two implementation methods. For example, to achieve a reasonably accurate fit to the Gaussian-like signal shown in Figure 5.2, the second derivative parameterisation method requires two more knots than the direct cubic spline interpolation method.

The FlexKnot signal model is then implemented in the REACH data analysis pipeline to test its functionality. In the cases of only a few knots, the FlexKnot function modelled via parameterising the second derivative is suitable for describing shapes with obvious troughs or crests, such as the shape of the global 21-cm signal, which, typically, is predicted to have one bigger trough and an ensuing emission between $z \sim 6 - 30$ (corresponding to $\sim 45 - 200$ MHz). To describe a smooth shape like the power law function accurately, it would require a much higher number of knots distributed throughout the argument range, as the second derivative of a power law is non-linear except when the order is 2 or 3. This characteristic can also potentially help separate the global 21-cm signal from the smooth foregrounds. Because the foregrounds are well-approximated by power laws, a FlexKnot model with a limited number of knots is structurally ill-suited to fitting them. This inherent stiffness against smooth, simple shapes means the model is less likely to accidentally absorb the foreground component, allowing it

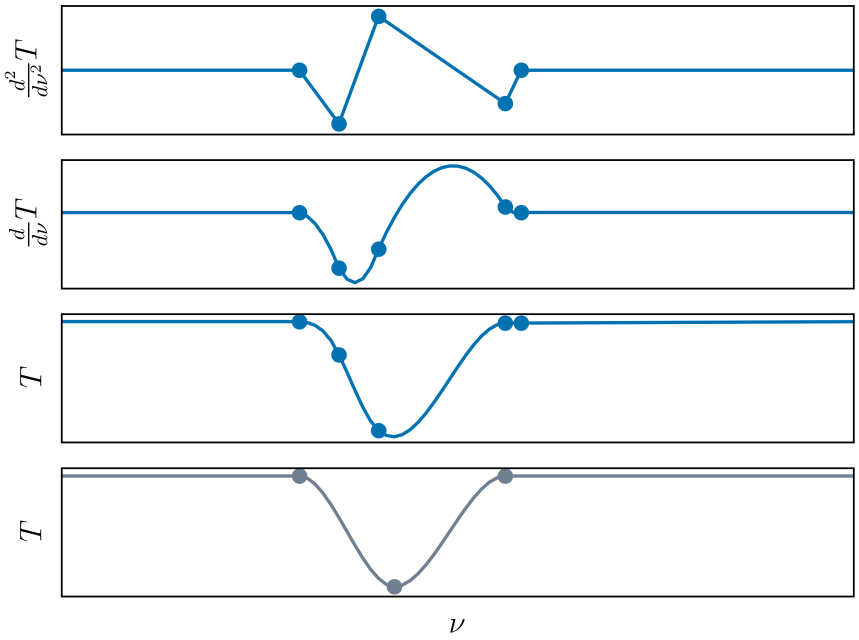


Figure 5.2: The upper three panels (blue) demonstrate the implementation of the FlexKnot function in three steps: the second derivative of the global 21-cm signal is parametrised by linearly interpolated splines, shown in the first panel, and the signal is recovered by integrating it twice, shown respectively in the second and the third panel. The bottom panel (gray) shows a similar signal directly fitted by Piecewise Cubic Hermite Interpolating Polynomial (PCHIP). Recovering a Gaussian-like signal requires a different number of knots for the two methods.

to preferentially capture the more complex, non-power-law signature of the global 21-cm signal. Figure 5.3 shows an example of how the FlexKnot model of different $N_{\rm knot}$ aims to recover a power law function:

$$T(\nu) = \nu^{-\frac{35}{2}}. \tag{5.2}$$

In this case, it requires 9 knots for it to be properly recovered, and a lower number of knots leads to inaccurate results. The optimal number of knots to use in each case would also depend on the step size in numerical integration.

5.2.1 Parameters and Priors

Each knot K_i is defined by two parameters: the position in the observing frequency and the second derivative in temperature. An additional highest frequency knot K_{N+1} fixed at 0 K, ensures the integrated value remains zero beyond this point, as it is known that reionisation will eventually eliminate the 21-cm signal. This

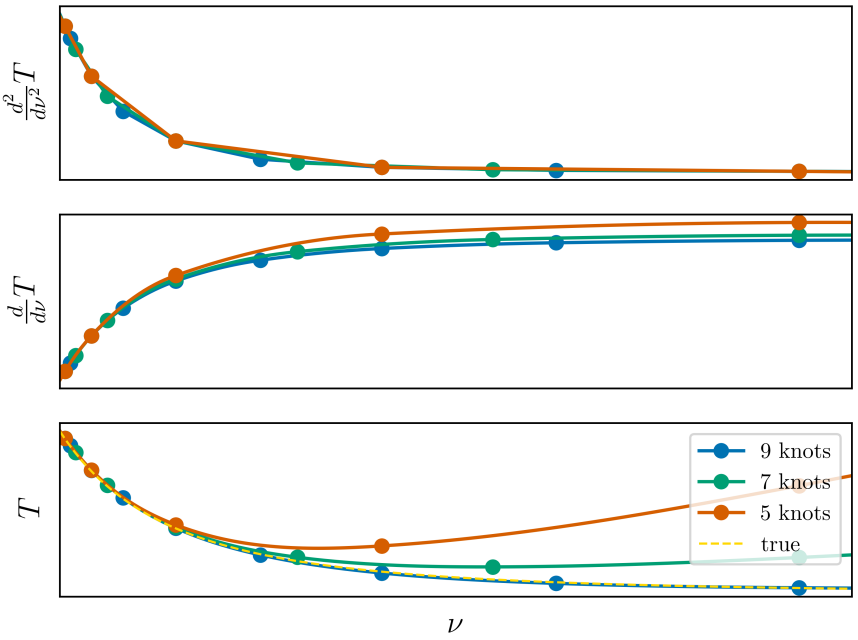


Figure 5.3: An illustration of the FlexKnot model's three-step process for recovering a power-law function, $T(\nu) = \nu^{-\frac{5}{2}}$ (yellow dashed line). **Top panel:** The model first parameterises the function's second derivative using linearly interpolated knots. **Middle panel:** This is then integrated once to recover the first derivative. **Bottom panel:** A second integration recovers the final function. The figure demonstrates that, unlike functions with a single peak or trough, an accurate recovery of a smooth power law requires a large number of knots distributed across the entire frequency range.

additional highest frequency knot is excluded from the total number of knots N_{knot}. The total number of parameters n_{signal} of the FlexKnot signal model is:

$$n_{\text{signal}} = 2 \times N_{\text{knot}} + 1. \tag{5.3}$$

As the FlexKnot model is implemented within a Bayesian framework, priors must be set for all parameters. For the positional parameters of the knots, a sorted prior is adopted. This ensures that the knot positions are always sampled in ascending order of frequency ($p_1 < p_2 < \cdots < p_{N_{\text{knot}}}$). The effect of this prior is illustrated in the corner plots in Figure 5.4, where the prior samples (blue dots) are confined to the upper triangle of the plot, enforcing the ordering. The prior range for these positional parameters is set to $[50, 200]$ MHz (sorted uniform prior), which covers the observational frequency range of the REACH experiment.

For the second derivative parameters, the prior is set to be $[-0.02, 0.02]$ (uniform prior); the prior limits should be adjusted depending on the interval between the frequency points, as the derivative describes the rate of change between two

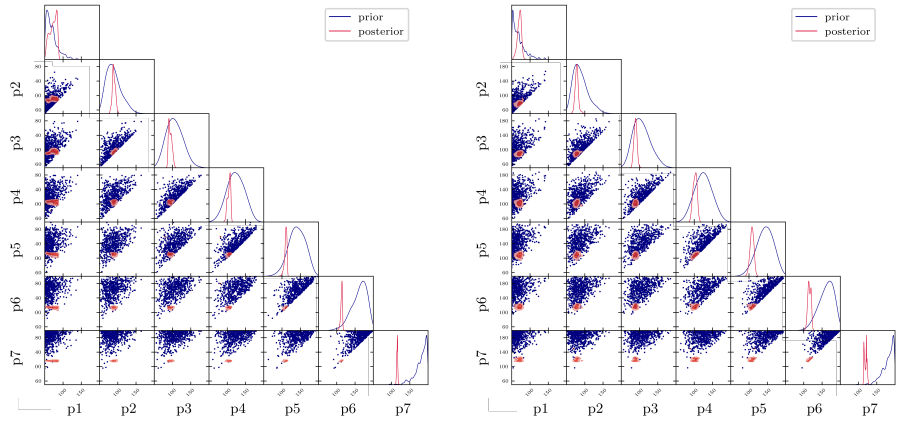


Figure 5.4: Corner plots showing the prior (blue) and posterior (red) distributions for the positional parameters (p_n) of the knots in a 6-knot FlexKnot model. The **left panel** shows a case without systematics, while the **right panel** includes a 20 mK sinusoidal systematic. The injected signal is a Gaussian centred at 85 MHz with an amplitude of 255 mK and a standard deviation of 15 MHz. The narrower posterior distributions in the left-hand case demonstrate that the knot positions are well constrained by the data in the absence of systematics.

discrete points in this model. In these tests, the interval between the frequency points is set to 1 MHz. The second derivative of the lowest frequency point (the first knot K_1), however, is set to be negative only $[-0.02, 0]$ (uniform prior) to prevent the function from turning uncharacteristically high at the low frequencies corresponding to higher redshifts, which could potentially occur when the chromatic foregrounds in the form of power law function are present in the data. The additional highest frequency knot K_{N+1} does not have the second derivation parameter; its twice integrated value is fixed at 0 K, with both its first and second derivative being also null.

5.2.2 Theoretical dark ages Primer

Unlike during the EoR, the theoretical global 21-cm signal is better understood during the dark ages (Hogan and Rees, 1979; Scott and Rees, 1990; Mondal and Barkana, 2023). To take advantage of that, in this model (Figure 5.5), a set of theoretical values at frequencies corresponding to the dark ages ($z \sim 30 - 40$, or $\sim 35 - 45$ MHz) lying outside the observational frequency range of REACH is implemented as a primer, or the first segment that initialises the ensuing function, to not only provide a reasonable and reliable constraint but also further prevent potential uncharacteristic positive surge at the lower frequencies that sometimes occurs, as the case shown in Figure 5.6. The low-frequency end of the function is then interpolated between the last point of the theoretical dark ages primer and the first knot. The theoretical dark ages signal values are generated using the

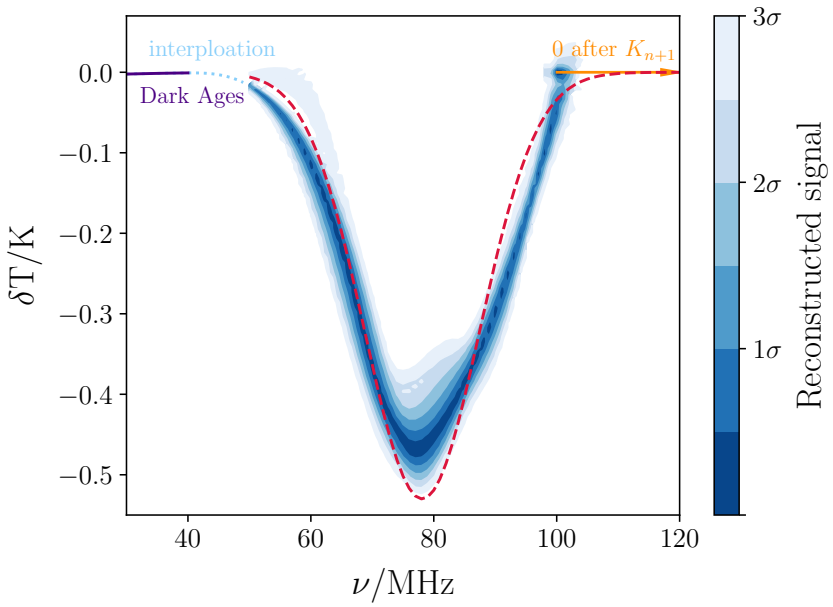


Figure 5.5: An example of a signal reconstruction using the 4-knot FlexKnot model. The injected signal in the shape of a Gaussian is shown as a red dashed line, while the reconstructed signal is represented by the blue contours. The model is constrained by: a theoretical dark ages primer between 30 MHz and 40 MHz (purple solid line), a negative second-derivative prior on the first knot, and a requirement that the signal returns to zero at high frequencies after the highest frequency knot K_{N+1} (orange arrow). The blue dotted line traces the interpolation between 40 MHz and 50 MHz, the endpoint of the dark ages primer and the beginning of the reconstructed signal.

GLOBALEMU emulator (Bevins et al., 2021), modified with parameters specific to the dark ages. The global 21-cm signal during the dark ages is characterised by the following astrophysical parameters: baryon density, $\Omega_b = 0.01495074$, matter density, $\Omega_m = 0.29395689$, curvature density, $\Omega_k = 0.05966342$, and dimensionless Hubble constant, $h = 0.75665809\,\mathrm{km\,s^{-1}Mpc^{-1}}$.

5.3 Simulated Antenna Temperature Data

The simulated antenna temperature data generated to test the capability of the FlexKnot signal model to extract the global 21-cm signal in the REACH data analysis pipeline are described in this section. Based on the established model introduced in Section 2.4.1, the data include the following four components: the foregrounds, the sky-averaged 21-cm signal, the systematics, and the antenna temperature noise:

$$T_{\mathrm{data}} = T_{\mathrm{f}} + T_{21} + T_{\mathrm{sys}} + T_{\mathrm{noise}}. \tag{5.4}$$

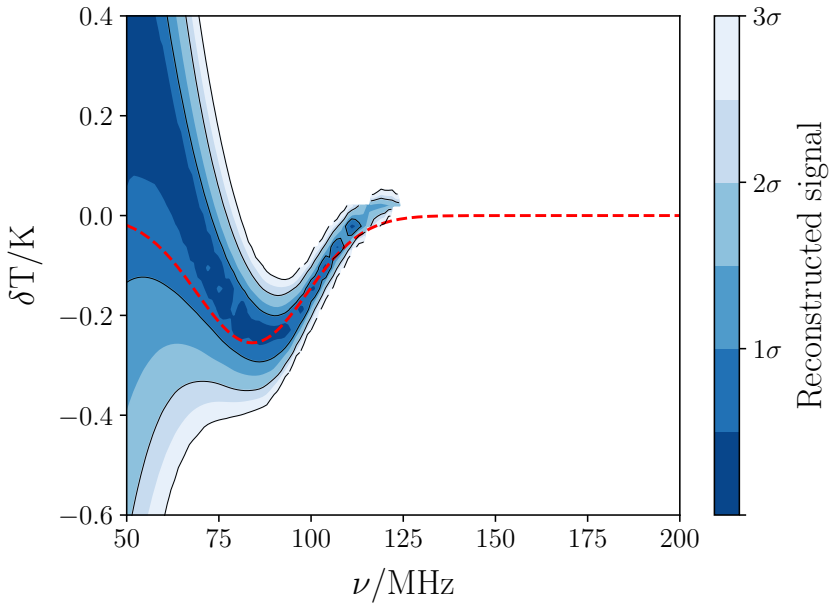


Figure 5.6: An example of a signal reconstruction using a 6-knot FlexKnot model without its low-frequency constraints; no dark ages primer and no negative prior on the first knot. The injected signal is shown by the red dashed line, and the reconstructed signal by the blue contours. Without the priors, the reconstructed signal diverges to unphysical positive values at low frequencies. However, features of the absorption trough, such as its centre frequency, depth, and width, are still reasonably recovered.

Foreground

The spatially varying spectral index map is simulated by the sky model based on instances of the 2008 Global Sky Model (de Oliveira-Costa et al., 2008) at 408 MHz and 230 MHz:

$$\beta(\theta,\phi) = \frac{\log\left(\frac{T_{230}(\theta,\phi)-T_\gamma}{T_{408}(\theta,\phi)-T_\gamma}\right)}{\log\left(\frac{230}{480}\right)}, \tag{5.5}$$

with which the sky model can be generated:

$$T_{\text{sky}}(\theta,\phi,\nu) = (T_{480}(\theta,\phi) - T_\gamma)\left(\frac{\nu}{408}\right)^{-\beta(\theta,\phi)} + T_\gamma. \tag{5.6}$$

It is then convolved with the beam pattern of a conical log spiral antenna (Dyson, 1965) of beam pattern $D(\theta,\phi,\nu)$ to generate the foreground component:

$$T_{\text{f}} = \int_\Omega D(\theta,\phi,\nu) T_{\text{sky}}(\theta,\phi,\nu)\,\mathrm{d}\Omega. \tag{5.7}$$

Sky-averaged 21-cm signal

The Gaussian signal, the flattened Gaussian signal, and the signals emulated by GLOBALEMU are the three different types of signal that have been injected as the

true signal in the simulated antenna temperature data in this work.

We shall take a Gaussian model of the global 21-cm signal as

$$T_{21,\mathrm{g}}(\nu) = -A_{21}\left(e^{\frac{-(\nu-\nu_0)^2}{2\sigma^2}}\right), \tag{5.8}$$

where A_{21} is the absorption amplitude, ν_0 is the centre frequency, and σ is the standard deviation. The explored parameter space of the injected signal in the form of a Gaussian is as follows:

- amplitude: $\{155, 255, 455\}$ mK
- centre frequency: $\{85, 105, 125\}$ MHz
- standard deviation: $\{10, 15, 20\}$ MHz

Bowman et al. (2018) defined the flattened Gaussian function with the equation:

$$T_{21,\mathrm{f}}(\nu) = -A_{21}\left(\frac{1-e^{-\tau e^{B}}}{1-e^{-\tau}}\right), \tag{5.9}$$

with

$$B = \frac{4(\nu-\nu_0)^2}{w^2}\log\left[-\frac{1}{\tau}\log\left(\frac{1+e^{-\tau}}{2}\right)\right], \tag{5.10}$$

where w is the full width at half maximum, and τ is a flattening factor. The unsmooth flat bottom of a flattened Gaussian is what distinguishes it from the two other types of signal. The parameters used to characterise the flattened Gaussian signal EDGES claimed to have detected have the following values: $A_{21} = 0.52$ K, $\nu_0 = 78.3$ MHz, $w = 20.7$ MHz, and $\tau = 6.5$.

GLOBALEMU (Bevins et al., 2021) is a sky-averaged 21-cm signal emulator, and the signal during cosmic dawn and the epoch of reionisation is characterised by the following astrophysical parameters: the star formation efficiency, f_*, the minimal virial circular velocity, V_c, the X-ray efficiency, f_X, the CMB optical depth, τ, the slope of the spectral energy density, α, the low energy cut-off of the X-ray spectral energy density, ν_min, and the mean free path of ionising photons, R_mfp.

Table 5.1 lists the parameters used to generate the signals adopted in this chapter. Some signals are emulated with an enhanced radio background, where the parameter α is replaced by f_radio, for the deeper absorption troughs that are easier to recover. To clarify the distinction between the two models, the parameter α in the standard model controls the efficiency of X-ray heating of the gas (affecting the spin temperature), while f_radio in the non-standard models parametrises a hypothetical excess radio background (affecting the radio background temperature). This physical difference is what allows the latter models to produce the deeper absorption troughs needed to explain the EDGES result (Sharma, 2018; Fialkov and Barkana, 2019), the detected signal of which has an amplitude incompatible with the standard 21-cm models. For the purpose of this work, only signals with absorption trough deeper than 15 mK are used as the true signal.

Table 5.1: A table of parameters used to generate the GLOBALEMU signals tested in this chapter. The first four cases are generated using the original version of GLOBALEMU, while the last two are generated with an enhanced radio background, where the parameter α is replaced by $f_{\rm radio}$. The subscript indicates the frequency corresponding to the deepest point of the signal's absorption trough.

	f_*	$V_{\rm c}$	$f_{\rm X}$	τ	α	$\nu_{\rm min}$	$R_{\rm mfp}$
G_{73}	2.746×10^{-1}	3.795×10^{1}	7.504×10^{1}	3.591×10^{-2}	1.242	2.914	3.262×10^{1}
G_{85}	1.715×10^{-1}	5.231×10^{1}	4.756×10^{2}	4.303×10^{-2}	1.188	2.920	2.297×10^{1}
G_{99}	3.657×10^{-1}	9.424×10^{1}	2.053×10^{2}	7.732×10^{-2}	1.078	2.814	4.573×10^{1}
G_{107}	4.176×10^{-1}	8.783×10^{1}	3.459×10^{1}	7.104×10^{-2}	1.032	2.562	3.842×10^{1}

	f_*	$V_{\rm c}$	$f_{\rm X}$	τ	$f_{\rm radio}$	$\nu_{\rm min}$	$R_{\rm mfp}$
G^*_{73}	4.113×10^{-2}	4.490×10^{1}	1.332×10^{0}	5.753×10^{-2}	2.112×10^{2}	0.000	4.000×10^{1}
G^*_{99}	1.588×10^{-1}	1.532×10^{1}	1.096×10^{0}	9.093×10^{-2}	1.129×10^{2}	0.000	4.000×10^{1}

Systematics

In the cases where systematics are present, the simulated antenna temperature data include a sinusoid with an amplitude of 20 mK, 50 mK, or 100 mK. The phase of the sinusoids is shifted with respect to the centre frequency of the injected signal or the frequency at which the signal is at its lowest, so the phase difference between the signal and the sinusoid is the same across all cases:

$$T_{\rm sys} = A_{\rm sys} \sin\left(2\pi \frac{(\nu - \nu_0)}{P_{\rm sys}} + \phi_{\rm sys}\right), \tag{5.11}$$

where $A_{\rm sys}$ is the amplitude, $P_{\rm sys}$ is the period, and $\phi_{\rm sys}$ is the phase.

Antenna temperature noise

The antenna temperature noise is generated by the generalised normal distribution noise model where the noise at each frequency channel is randomly assigned based on the probability density distribution. It is a reasonable alternative to the radiometric noise and is easier to generate. The scale parameter of the generalised normal distribution noise model can be expressed as

$$\alpha_{\rm n} = \sigma_{\rm n} \sqrt{\frac{\Gamma\left(\frac{1}{\beta_{\rm n}}\right)}{\Gamma\left(\frac{3}{\beta_{\rm n}}\right)}}, \tag{5.12}$$

where $\beta_{\rm n}$ is the shape parameter, $\sigma_{\rm n}$ is the standard deviation, and Γ denotes the gamma function. In this work, the shape parameter is set to $\beta_{\rm n} = 2$, and the standard deviation is $\sigma = 25$ mK unless otherwise specified.

5.4 Results

5.4.1 Model Comparison

In theory, the FlexKnot signal model should have the versatility to fit injected signals of different types of shapes, and for the cubic spline twice integrated from the linear spline method, shapes that are smooth without rigid turns are favoured. Signals of different types of shapes, namely, the Gaussian signal, the flattened Gaussian signal, and the GLOBALEMU signal, are injected in the data to test the model's capabilities and make comparison with the commonly used Gaussian signal model. Apart from $\Delta \log \mathcal{Z}$ (Equation 2.19), the RMSE between the injected signal, T_{21}, and the reconstructed signal, T_{21}^{*}, is also calculated to compare between different models (Equation 2.21).

Gaussian

The first test evaluates the FlexKnot model's performance against data containing a simple Gaussian signal. The parameter space for these signals is listed in Section 5.3. As the true signal shape perfectly matches the Gaussian model, Bayesian principles predict that this simpler, three-parameter model should be favoured over the more complex FlexKnot model.

The results, summarised in Table 5.2, confirm this expectation but also demonstrate the robustness of the FlexKnot model. While the Gaussian model consistently yields a slightly lower signal RMSE and higher log evidence, the FlexKnot model nevertheless recovers the signal with high statistical significance ($\Delta \log \mathcal{Z}$) in most cases. Figure 5.7 shows a representative example, plotting the metrics for both the FlexKnot model in solid lines and the reference Gaussian model in dashed lines.

The figure also illustrates the impact of instrumental systematics on the recovery. For both models, the signal RMSE increases progressively with the amplitude of the added systematic, with the recovery becoming unreliable for systematics of 50 mK or greater. This effect is discussed in more detail in Section 5.4.2.

As shown in Table 5.2 and Figure 5.7, the dedicated Gaussian model and its variant yield marginally lower signal RMSE's and higher Bayesian ratios. The additional parameters of the FlexKnot model are penalised when they are not required to fit the data. This result serves as an important validation, demonstrating that FlexKnot does not over-fit simple signals and that the Bayesian evidence correctly favours the simpler model.

GLOBALEMU

The merit of the FlexKnot model is demonstrated in tests using physically motivated signals from the GLOBALEMU emulator. As shown in Table 5.3, it consistently

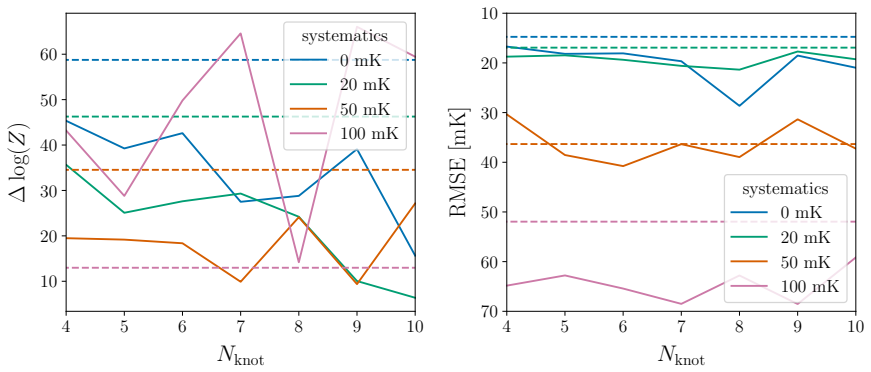


Figure 5.7: Comparison of the FlexKnot (solid lines) and Gaussian (dashed lines) signal models when recovering an injected Gaussian signal. The **left panel** shows the log Bayes ratio, $\Delta \log \mathcal{Z}$, while the **right panel** shows the signal RMSE as a function of the number of knots, N_{knot}, for different levels of added sinusoidal systematics (different colours). The injected signal is centred at 85 MHz with an amplitude of 155 mK and a standard deviation of 10 MHz. Note that the RMSE axis is inverted. The plot shows that while the FlexKnot model achieves confident detections, the simpler Gaussian model yields slightly lower RMSE's. The signal RMSE panel also shows that the quality of the fit degrades as the amplitude of the systematic increases. The signal recovery becomes unreliable for systematic amplitudes of 50 mK and 100 mK, which yield significantly higher RMSE values.

outperforms the Gaussian model, proving flexible enough to accurately recover the diverse absorption histories produced by different sets of astrophysical parameters (listed in Table 5.1).

A representative result is shown in Figure 5.8. The high $\Delta \log \mathcal{Z}$ indicates a confident detection in all cases, and the signal RMSE confirms that the FlexKnot model yields a significantly more accurate signal recovery than the Gaussian model.

While the FlexKnot model's performance is superior, its signal RMSE may still appear high relative to the ~ 145 mK signal depth. This is due to the continuous significant emission at the higher frequencies in this particular case. The FlexKnot model is constrained to return to zero after the final knot and is therefore unable to fit this emission. This is shown in Figure 5.9. This discrepancy at high frequencies contributes significantly to the total RMSE; if the emission range is excluded from the calculation, the RMSE of the fit to the absorption trough is much lower.

This test demonstrates the primary strength of FlexKnot: its ability to accurately recover physically plausible, non-Gaussian signal shapes that would be mischaracterised by simpler models. The signals generated by GLOBALEMU often exhibit asymmetries and emission features that cannot be captured by a simple Gaussian. As shown in Table 5.3 and Figure 5.8, the FlexKnot model consistently yields a lower signal RMSE than the Gaussian model for these signals. While the

Table 5.2: Signal RMSE and $\Delta \log \mathcal{Z}$ comparison between the FlexKnot model $\mathcal{M}_{\rm FK}$ and the Gaussian model $\mathcal{M}_{\rm G}$. The injected signal is in the shape of a Gaussian. $\mathcal{M}_{\rm G}$ yields a lower signal RMSE in most of the cases and a higher $\Delta \log \mathcal{Z}$ in all of the listed cases.

Signal parameter			**RMSE**		**$\Delta \log \mathcal{Z}$**	
ν_0	σ	A_{21}	$\mathcal{M}_{\rm FK}$	$\mathcal{M}_{\rm G}$	$\mathcal{M}_{\rm FK}$	$\mathcal{M}_{\rm G}$
85	10	155	16.7	14.8	45.3	58.7
85	15	155	15.8	15.9	41.0	61.1
85	20	155	22.6	19.8	39.7	52.0
105	10	155	25.1	23.9	78.6	94.0
105	15	155	18.9	32.2	78.6	94.9
105	20	155	10.5	14.2	68.7	84.3
125	10	155	14.6	2.3	85.7	91.7
125	15	155	11.8	7.9	103.5	114.5
125	20	155	8.0	3.0	107.2	119.7
85	10	255	21.3	11.7	90.0	96.0
85	15	255	20.8	25.4	87.1	95.0
85	20	255	51.9	34.2	65.9	85.8
105	10	255	30.9	11.3	127.0	145.6
105	15	255	24.6	5.7	133.2	142.9
105	20	255	10.4	7.3	128.3	137.2
125	10	255	14.6	3.6	143.3	154.4
125	15	255	19.6	5.3	170.6	177.6
125	20	255	18.7	5.4	176.4	184.6
85	10	455	24.6	14.3	158.4	164.7
85	15	455	27.4	18.6	142.0	162.2
85	20	455	44.4	26.0	129.1	146.8
105	10	455	49.0	21.1	220.5	239.4
105	15	455	19.2	28.4	225.4	244.2
105	20	455	17.6	12.7	217.2	232.7
125	10	455	17.1	2.5	245.4	261.5
125	15	455	29.4	4.4	306.8	313.6
125	20	455	21.3	3.1	330.3	346.0

Gaussian model can approximate the main absorption trough, it fails to capture the full morphology, which would lead to biased estimates of the underlying astrophysical parameters. FlexKnot's ability to trace the true signal shape makes it a far more powerful tool for physical interpretation.

Flattened Gaussian

The flattened Gaussian profile, motivated by the detection claim from Bowman et al. (2018), provides a unique test case. Its defining flat bottom and sharp corners make the overall shape resemble a piecewise linear spline more than a piecewise cubic spline.

When fitting this signal, the Gaussian model is, by its nature, unable to accurately capture the profile's shape, although it can recover the approximate centre frequency. The FlexKnot model performs significantly better, successfully cap-

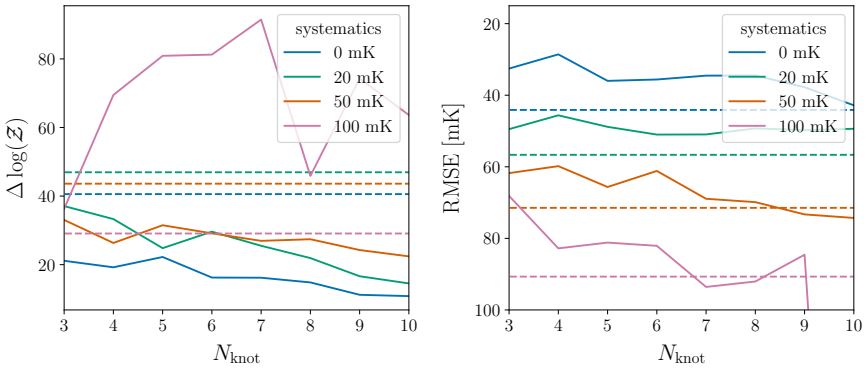


Figure 5.8: Comparison of the FlexKnot (solid lines) and Gaussian (dashed lines) signal models when recovering a GLOBALEMU signal. The left panel shows the log Bayes ratio, $\Delta \log \mathcal{Z}$), and the right panel shows the signal RMSE as a function of the number of knots, N_{knot}. Different colours represent varying levels of added sinusoidal systematics. The injected signal (G_{73} from Table 5.1) has a complex, non-Gaussian shape. Note that the RMSE axis is inverted. The figure demonstrates that the FlexKnot model consistently yields a lower RMSE, with an optimal number of knots between 4 and 6. The decreasing $\Delta \log \mathcal{Z}$ indicates that adding more knots does not improve the confidence in the signal recovery. Furthermore, similar to the results shown in Figure 5.7, the quality of the fit degrades as the amplitude of the systematic increases. The signal recovery becomes unreliable for systematic amplitudes of 50 mK and 100 mK.

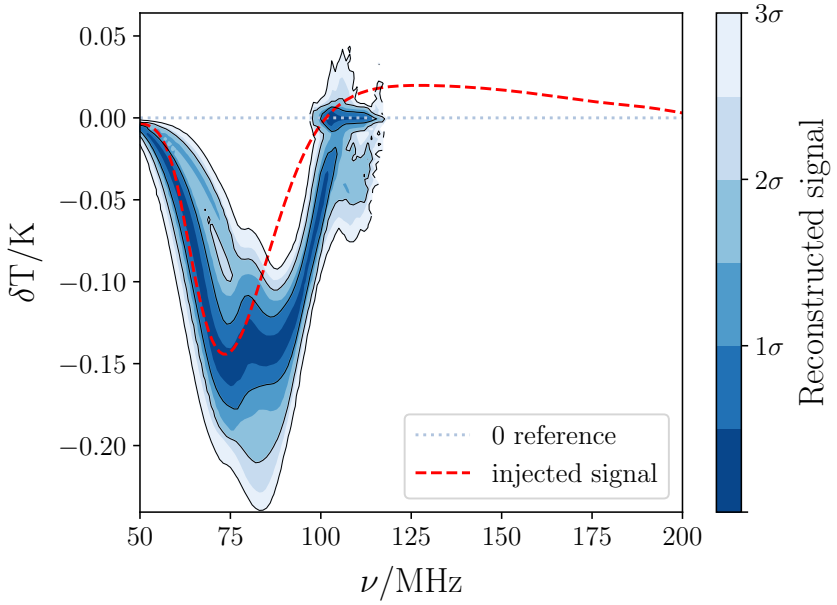


Figure 5.9: An example of a 4-knot FlexKnot signal reconstruction (blue contours) for a GLOBALEMU signal (G_{73} in Table 5.1) that exhibits a significant emission feature at high frequencies (red dashed line). The plot demonstrates a constraint of the model: because the signal is forced to zero after the final knot, it is unable to capture the high-frequency emission feature.

Table 5.3: Comparison of the signal RMSE and $\Delta \log \mathcal{Z}$ for the FlexKnot $\mathcal{M}_{\rm FK}$ and Gaussian $\mathcal{M}_{\rm G}$ models when fitting injected signals generated by GLOBALEMU. The FlexKnot model consistently yields a lower signal RMSE in all cases. The subscript indicates the frequency corresponding to the deepest point of the signal's absorption trough.

	RMSE		$\boldsymbol{\Delta \log \mathcal{Z}}$	
	$\mathcal{M}_{\rm FK}$	$\mathcal{M}_{\rm G}$	$\mathcal{M}_{\rm FK}$	$\mathcal{M}_{\rm G}$
G_{73}	28.6	44.1	22.2	40.6
G_{85}	42.9	46.9	24.9	43.5
G_{99}	17.4	25.7	53.9	69.6
G_{107}	21.4	35.5	71.5	93.1

turing the signal's overall width and location. However, it also faces a challenge with the sharp corners. Because the FlexKnot model is constructed from smooth, twice-integrated splines, it is structurally ill-suited to reproducing sharp changes in curvature and tends to round off the edges of the trough, as illustrated in Figure 5.10.

In the presence of chromatic foregrounds and noise, the model often fails to capture the sharp features of the flattened Gaussian. Instead of allocating the 4 knots required to form a flat bottom, the analysis typically allocates fewer knots, resulting in a smoother, more rounded shape. This is demonstrated in the example shown in the bottom of Figure 5.10: even with 7 knots available, the model fails to correctly place the four necessary to resolve the flat feature.

This behaviour is a result of Bayesian model selection. The analysis penalises the high model complexity required to force the smooth splines to fit sharp corners, particularly when those features are obscured by a high noise level (in this case, $\sigma = 25$ mK). It consistently finds a simpler solution in a smoother trough. This also explains why simply increasing the number of available knots is not a solution; the model's extra flexibility begins to erroneously fit foreground residuals rather than the signal's true shape. The recovery of such sharp features could be improved by either lowering the noise level in the data or by adopting a different model construction, such as one based on piecewise linear splines.

The results show that while the model captures the overall location and width of the trough, its inherent smoothness struggles to replicate the sharp spectral breaks of the flattened Gaussian (Figure 5.10). In the presence of foregrounds and noise, the model preferentially fits a smoother trough, often with fewer knots than required to form the flat bottom. This finding suggests that distinguishing an intrinsically sharp spectral feature from noise or unmodelled systematics is exceptionally difficult without a highly specific, targeted model. It highlights the challenge any physics-agnostic model faces when confronted with signal shapes possessing sharp derivatives.

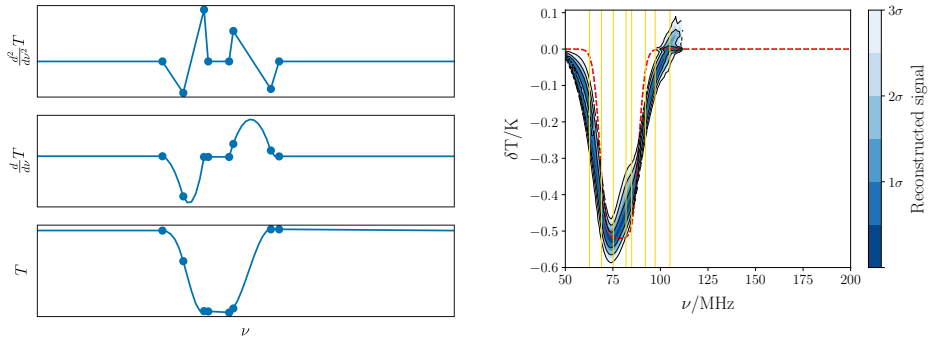


Figure 5.10: An illustration of the FlexKnot model's difficulty in reproducing the sharp features of a flattened Gaussian signal. **Left:** An idealised three-step construction showing how the model parameterises the second derivative (top), first derivative (middle), and final signal (bottom). This demonstrates that a minimum of four knots are required to form the flat bottom of the trough. **Right:** An example of an actual signal reconstruction from a simulation where the injected signal is a flattened Gaussian. The yellow vertical lines indicate the mean posterior positions of the seven available knots. The plot shows that the model fails to allocate the four knots necessary to resolve the flat feature, instead finding a simpler, smoother solution with three knots. This behaviour persists even when a higher number of knots is used.

5.4.2 Signal Recovery with Systematics

To explore the FlexKnot model's robustness against instrumental systematics, a series of tests were conducted using simulated data containing sinusoidal contaminants of varying amplitudes (0 mK, 20 mK, 50 mK, and 100 mK). Each case is tested using a range of N_{knots}, of which the lowest signal RMSE is recorded. The quality of the fit is evaluated using the signal RMSE, normalised by the true signal's amplitude to yield a dimensionless value. The results are summarised in Figure 5.11, where the cases are grouped by the amplitude, the centre frequency, and the standard deviation of the injected signal.

As expected, a higher level of systematic contamination consistently leads to a larger signal RMSE, but the degree of this degradation depends on the signal's properties. The recovery is more robust for signals with a deeper absorption trough (higher amplitude), as they have a greater signal-to-contaminant ratio. The signal's centre frequency also has a minor effect; in the absence of systematics, signals at higher frequencies are recovered slightly better, likely because the intrinsic foreground is less bright. This advantage disappears in the presence of systematics, which, depending on the its phase, can potentially distort the signal profile. No significant dependence on the signal's standard deviation was observed, likely because the period of the injected sinusoid was larger than the signal widths tested ($10 - 20$ MHz).

Overall, the results show that when antenna temperature data contain sinus-

oidal systematics with an amplitude higher than 50 mK, the injected 21-cm signal cannot be recovered with sufficient accuracy using the FlexKnot signal model.

5.4.3 Optimal Number of Knots

This section investigates the optimal number of knots that could represent the global 21-cm signal. This is done by running the algorithm using a range of number of knots, $N_{\text{knot}} = 4 - 10$, and examining the fitted signals by log Bayes ratio and signal RMSE.

Data without systematics

The optimal number of knots required to model both Gaussian and GLOBALEMU signals is consistently found to be between 4 and 6. This conclusion is supported by multiple diagnostics: the log Bayes ratio, $\Delta \log \mathcal{Z}$, generally peaks in this range, while the signal RMSE is minimised, as shown in Figure 5.12.

Adding knots beyond this optimal range does not improve the signal recovery. Instead, it often degrades the fit by introducing unwanted structures, as the model's extra degrees of freedom begin to erroneously fit noise or foreground residuals. Interestingly, this finding aligns with the results of Heimersheim et al. (2023), who, despite using a different implementation, also found that the Bayesian evidence peaked at $N_{\text{knot}} = 6$ using EDGES low-band data.

The optimal number of knots for a flattened Gaussian signal, however, depends strongly on the presence of foregrounds and noise. In an ideal, noise-free scenario, approximately $7-8$ knots are required to resolve the sharp corners and flat bottom of the trough.

In a more realistic case where the signal is combined with chromatic foregrounds and noise, these sharp features become masked. The analysis then regresses the signal fit to a simpler model with a smoother, rounded bottom, the degree of which depends on the noise level. Consequently, the log evidence peaks at a lower number of knots ($N_{\text{knot}} \sim 5 - 6$), similar to the number required for a standard Gaussian or GLOBALEMU signal.

Data with systematics

In the presence of lower-level systematics ($A_{\text{sys}} < 50$ mK), where the signal can still be recovered with high confidence, the optimal number of knots remains unchanged at $N_{\text{knots}} \sim 4 - 6$. Adding knots beyond this range does not improve the signal recovery; instead, the model's extra degrees of freedom begin to erroneously fit the systematic and foreground residuals. $\Delta \log \mathcal{Z}$ does not improve with the increasing N_{knots} either.

In the high-systematics regime ($A_{\text{sys}} > 50$ mK), a different behaviour is observed. Here, increasing the number of knots can cause a significant but misleading

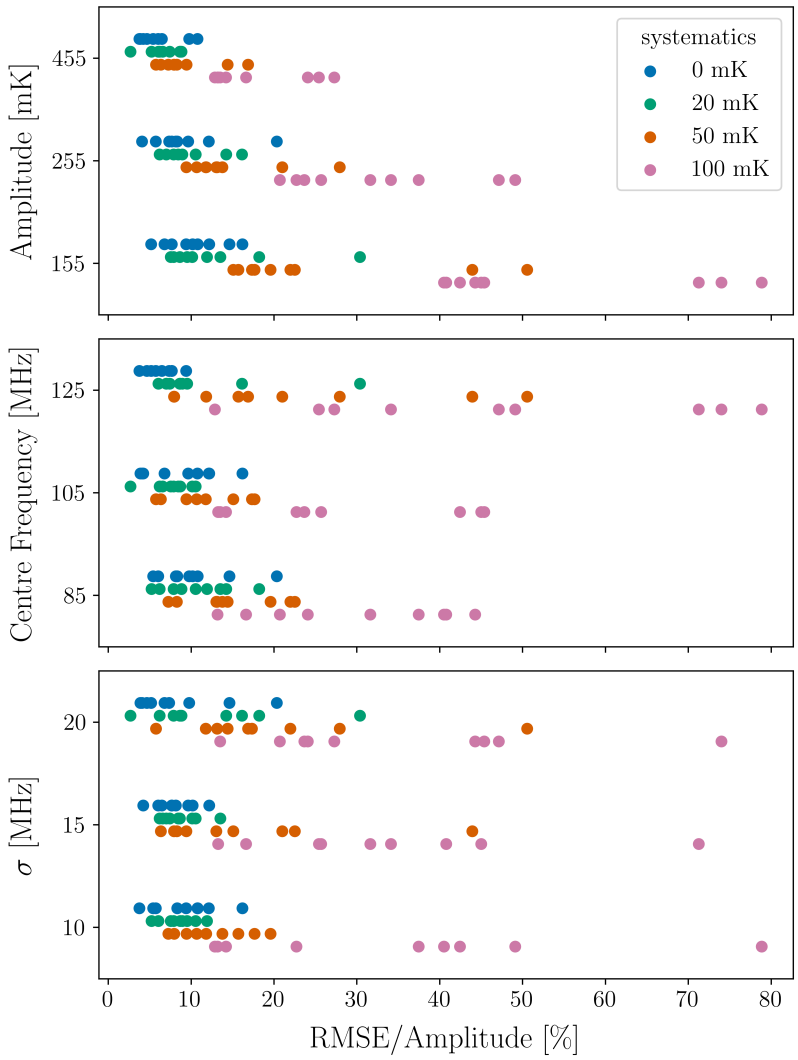


Figure 5.11: The normalised signal RMSE (signal RMSE divided by the injected signal's amplitude) as a function of injected Gaussian signal parameters and the level of added systematics (indicated by colour). The three panels present a full 3D parameter sweep marginalised onto the signal's amplitude (top), centre frequency (middle), and standard deviation (bottom) from the the same set of data. The **top** panel shows that signal recovery is generally more robust for signals with a deeper absorption level. The **middle** panel indicates that while signals at higher centre frequencies are recovered slightly better in the absence of systematics, this advantage disappears when systematics are introduced. The **bottom** panel reveals no significant dependence on the signal's standard deviation. Overall, the signal recovery becomes unreliable for all cases once the systematic's amplitude reaches 50 mK.

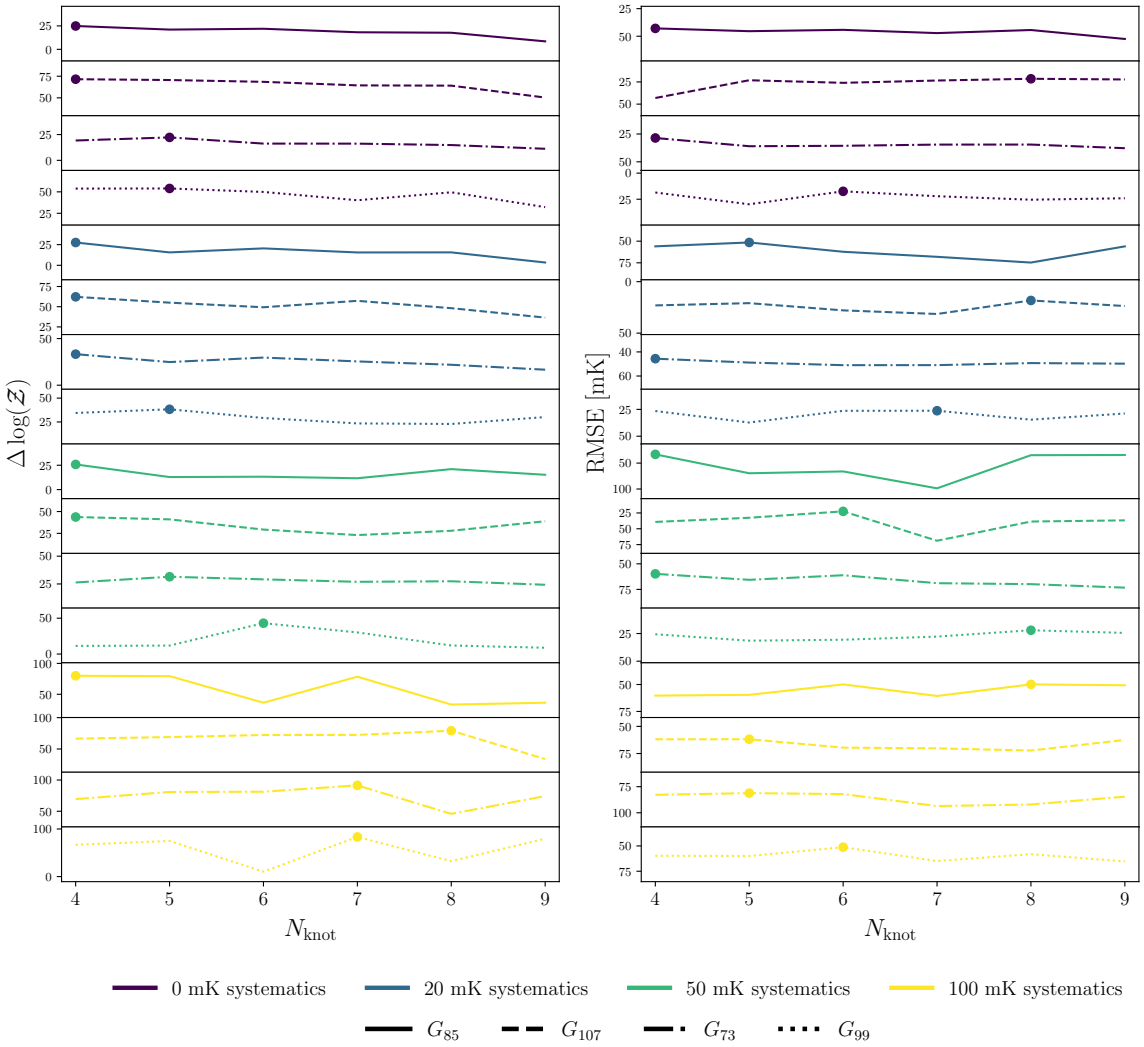


Figure 5.12: The dependence of the log Bayes ratio ($\Delta \log \mathcal{Z}$, left) and the signal RMSE (right) on the number of knots. Different line styles correspond to different injected GLOBALEMU signals (see Table 5.1), while the colours indicate the amplitude of an added sinusoidal systematic. The optimal point for each case (highest $\Delta \log \mathcal{Z}$ or lowest RMSE) is marked with a dot. Note that the y-axis for the RMSE is inverted, so a peak corresponds to a minimum error. The results consistently show that the optimal number of knots is in the range of $N_{\text{knot}} \sim 4 - 6$. Using more knots does not improve the signal recovery and often degrades the fit.

increase in the log Bayes ratio, $\Delta \log \mathcal{Z}$. This is, again, a result of the FlexKnot signal model starting to fit the sinusoidal systematics alongside the foreground residuals instead of the true signal; it is confidently fitting the other structures such as part of the foregrounds or the systematics other than the true signal in the given data.

5.5 Conclusions

In this work, it has been demonstrated that the physics-agnostic FlexKnot signal model, when integrated into a robust Bayesian pipeline, is a powerful and versatile tool for extracting the global 21-cm signal. The capability and performance of the FlexKnot signal model, implemented in the REACH data analysis pipeline, have been investigated for the recovery of the global 21-cm signal from simulated data, both with and without the presence of systematics and for several different injected signal implementations.

This FlexKnot signal model is characterised by a cubic function integrated twice from a set of freely moving knots interpolated by piecewise linear splines, a different implementation of the FlexKnot signal model from previous works (e.g. Heimersheim et al. (2023)). It is conditioned by the dark ages primer outside the REACH observing frequency range ($50 - 200$ MHz, which covers the period between cosmic dawn and reionisation) at the lower frequency end. The prior on the second derivative of the first knot (at the lowest frequency) is constrained to be negative. This prevents the signal model from producing a large, unphysical emission feature at high redshifts, which would contradict established theoretical predictions. An additional highest-frequency knot is also included to ensure the ensuing values at higher frequencies go to zero.

The FlexKnot signal model is implemented in the REACH data analysis pipeline, where the foreground and the signal are jointly fitted, to test how well the signal model manages to recover the injected signal. For comparison, the same tests are performed using the Gaussian signal model. Three types of global 21-cm signal have been used as the true signal: the Gaussian signal, the flattened Gaussian signal, and physically motivated signals emulated by GLOBALEMU. The FlexKnot signal model can recover all tested signals with confidence and with reasonably low signal RMSE and outperforms the Gaussian signal model when the true signal is not in the form of a Gaussian. However, it struggles to resolve narrow features that are not smooth, such as the flat bottom of a flattened Gaussian.

In the presence of sinusoidal systematics as a proxy for unaccounted-for systematics e.g. cable reflections, the FlexKnot signal model has the tendency to start fitting the systematics, as the number of knots increases. This can be avoided by setting fewer number of knots. For the tested signals, which all have absorption levels higher than 120 mK, the FlexKnot signal model can reliably recover the



signal with confidence and reasonably low signal RMSE when the sinusoid as systematic has an amplitude lower than 50 mK. The deeper the absorption trough is, the less susceptible it is to the presence of systematics. The accuracy in finding the centre frequency, or the minimum point, of the true signal is also affected by sinusoidal systematics.

The optimal number of knots to recover the absorption history of the physical global signal is $N_{\text{knot}} \sim 4 - 6$. When there are more than sufficient knots, the FlexKnot model begins to fit unwanted structures, such as foreground residuals, which yields poorer fits. Heimersheim et al. (2023) arrived at a similar conclusion, showing that the evidence peaks at $N_{\text{knot}} = 6$ using EDGES low-band data despite the different implementations and interpolation methods. It is about the same for a signal in the form of a flattened Gaussian, as the FlexKnot signal model often struggles to fit the flat bottom and eventually recovers a signal similar to the other two types of signal that are smooth throughout. The optimal number of knots remains the same even in the presence of sinusoidal systematics, and adding more knots would result in it starting to fit other unwanted structures such as the systematics and the foregrounds.

The primary strength of FlexKnot is its ability to avoid model-dependent bias, a critical issue confronting the field. The results show that while a simple Gaussian model fails to capture the complex shapes predicted by physical emulators like GLOBALEMU, FlexKnot can accurately recover these diverse absorption and emission histories. By enabling a more faithful inference of the underlying astrophysical parameters, this approach is directly applicable to interpreting and comparing the results from other instruments like EDGES and SARAS.

The FlexKnot model represents a valuable approach for the field; it allows the data to inform the signal shape, enabling a robust characterisation of a potential detection or, the placement of model-independent constraints in the case of a null result. For future experiments aiming to confirm or refute existing claims, a flexible and well-constrained approach like the one demonstrated here is an essential component for achieving a credible result.

Part IV

Conclusions

Chapter 6

Conclusions

O, now, for ever farewell the
tranquil mind! Farewell content!

W. Shakespeare

6.1 Conclusions

In this thesis, several models are developed to address the major challenges in the global 21-cm experiments, from time-varying chromatic effects to fundamental uncertainty in the global 21-cm signal itself. These investigations, conducted within the robust Bayesian framework of the REACH data analysis pipeline, suggest a path forward that moves towards analysis frameworks that embrace complexity and uncertainty in all aspects of the measurement.

Chapter 2 introduces the core statistical methods and numerical analysis techniques used throughout the thesis. It begins with an overview of Bayesian inference, with particular emphasis on nested sampling. This is followed by a description of the REACH data analysis pipeline, which is structured around physically motivated models.

Chapter 3 presents the morphology of the ionosphere and develops a corresponding model, focusing on the D and F layers as the primary sources of chromatic effects relevant to the global 21-cm signal. This model is then used to investigate the time variability of the ionosphere and its impact on foreground contamination. The analysis finds that not only do ionospheric effects fail to average out, but they also demand high-fidelity modelling. Specifically, it has been shown that

uncertainty in the ionospheric parameters should be constrained to below 5% to prevent foreground residuals from overwhelming a potential detection.

Chapter 4 examines contamination from Galactic linear polarisation induced by Faraday rotation. The underlying physics of polarisation is outlined, followed by the presentation of three different models of Galactic linear polarisation built to stress-test the time-separated REACH data analysis pipeline. The results demonstrate that the REACH pipeline significantly outperforms standard methods, which generally lack robustness against contamination from linearly polarised foregrounds. In most cases, it successfully recovers injected global 21-cm signals with an amplitude of 0.157 K, centred between 80 MHz and 120 MHz, even in the presence of unaccounted-for polarisation contamination.

The work on the ionosphere and Galactic polarisation demonstrates that these are not static contaminants but dynamic, chromatic systems whose effects do not average out over time. Most existing analysis methods are insufficient to handle this complexity. While the REACH pipeline shows significant robustness, the field as a whole must move towards a new standard for mitigating these contaminants.

Chapter 5 introduces the FlexKnot signal model, which is characterised by a twice-integrated cubic function defined through freely positioned knots interpolated by piecewise linear splines. This model is designed to recover features in the global 21-cm signal that may arise from previously unaccounted-for physics. It is demonstrated that the FlexKnot model offers a flexible and effective alternative to traditional models. For complex, physically-motivated signals, it is able to significantly reduce the signal recovery error compared to a Gaussian model, with an optimal configuration of just $4-6$ knots proving sufficient to capture a wide range of signal morphologies. In the framework of the REACH data analysis pipeline, it is capable of accurately recovering the signal from the foreground and features consistent with theoretical predictions.

The main message of this thesis is that the components of a global 21-cm experiment, the instrument, the systematics, and the signal, are deeply intertwined. The path forward is to co-design them with holistic analysis pipelines that model these complex interactions from the ground up. The analysis is a core part of the measurement itself. By embracing this complexity, we can pave the way for a definitive, credible discovery of the Cosmic Dawn and Epoch of Reionisation.

6.2 Future

Further progress in global 21-cm cosmology will depend on a coordinated, multifaceted strategy. Achieving a reliable detection will likely require concerted efforts on three fronts: refining physical models of systematics and chromatic effects, developing advanced techniques for signal separation, and establishing rigorous, community-wide validation frameworks.

Improving Physical Models

The ideal approach to addressing chromatic effects is through more sophisticated and data-driven physical models.

For ionospheric effects, direct measurement is probably required. While monitoring stations provide crucial data, future models must account for the ionosphere's spatial variability by incorporating measurements from multiple stations, each corresponding to the portion of the sky they observe. This method depends on the availability and completeness of ionospheric data, where the absence of D-layer measurements remains a significant limitation. Further research is needed to determine the optimal sky division strategy that balances spatial resolution with analytical robustness.

Similarly, advancements are needed for polarised foregrounds. A more advanced approach, Faraday synthesis, combines aperture synthesis, rotation measure synthesis imaging, and deconvolution into a single algorithm. As demonstrated by Bell, M. R. and Enßlin, T. A. (2012), Faraday synthesis can yield higher spatial resolution and dynamic range, representing a more promising direction for accurately reconstructing polarised intensity as a function of both Faraday depth and sky position.

Exploring Alternative Separation Techniques

Beyond direct physical modelling, advanced signal processing offers a complementary path for isolating the global 21-cm signal. Chromatic ionospheric distortions, for example, can introduce statistical variability over time into what would otherwise be a deterministic foreground. The wavelet scattering transform, known for its stability against such distortions, may serve as an effective tool. Since the global 21-cm signal is static, while ionospheric fluctuations are time-variable, their distinct statistical characteristics could be exploited for a clean separation.

Non-Parametric Modelling of Unknown Systematics

The FlexKnot model also offers avenues for future development. A significant extension would be to improve the framework by using FlexKnot for modelling unknown systematics non-parametrically. In parallel, it would be beneficial to evaluate the performance of FlexKnot using a wider range of theoretical global 21-cm signals, including those representing more unconventional or extreme scenarios. Such investigations should also incorporate a systematic bias analysis to quantify how model flexibility and knot placement may influence the recovered signal. One strategy could involve implementing two distinct FlexKnot components simultaneously: one, constrained by informative priors, to capture the broad and smooth global 21-cm signal; and a second, designed to represent sharper,



potentially oscillatory instrumental features. Bayesian methods could then quantitatively determine whether the data are best explained by a signal, a systematic, or a combination of both. This approach would not only enable the discovery and characterisation of unknown systematic structures beyond predefined functional forms but also provide a framework for understanding and mitigating potential reconstruction biases.

Community-Wide Validation

The validation of these advanced models and techniques requires a coordinated community effort. The most effective way to build confidence and compare different analysis pipelines is through blind mock data challenges. These challenges should feature simulated data with realistic, complex foregrounds and a variety of injected signal shapes; some physically motivated, others representing potential systematics. This would allow for a rigorous, unbiased comparison of our collective analysis capabilities and provide a clear path toward a credible detection.

References

[1] E. Shen, D. Anstey, E. de Lera Acedo, and A. Fialkov. Bayesian data analysis for sky-averaged 21-cm experiments in the presence of ionospheric effects. MNRAS, 515(3):4565–4573, Sept. 2022.

[2] E. Shen, D. Anstey, E. de Lera Acedo, and A. Fialkov. FlexKnot as a generalized model of the sky-averaged 21-cm signal at z 6-30 in the presence of systematics. MNRAS, 529(2):1642–1653, Apr. 2024.

[3] E. Shen, D. Anstey, M. Spinelli, et al. Bayesian data analysis for sky-averaged 21-cm experiments with contamination from linearly polarized foregrounds. MNRAS, 539(4):3336–3346, June 2025.

[4] V. Gladkova. Arithmetic regularity lemmas and applications. PhD thesis, Apollo - University of Cambridge Repository, 2024.

[5] K. I. Kellermann, I. I. K. Pauliny-Toth, and P. J. S. Williams. The Spectra of Radio Sources in the Revised 3c Catalogue. ApJ, 157:1, July 1969.

[6] R. Laureijs, J. Amiaux, S. Arduini, et al. Euclid Definition Study Report. arXiv e-prints, arXiv:1110.3193, Oct. 2011.

[7] A. Dey, D. J. Schlegel, D. Lang, et al. Overview of the DESI Legacy Imaging Surveys. AJ, 157(5):168, May 2019.

[8] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. ApJ, 142:419–421, July 1965.

[9] B. G. Keating, P. A. R. Ade, J. J. Bock, et al. BICEP: a large angular scale CMB polarimeter. In S. Fineschi, editor, Polarimetry in Astronomy, volume 4843 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pages 284–295, Feb. 2003. doi: 10.1117/12.459274.

[10] A. Kosowsky. The Atacama Cosmology Telescope. New A Rev., 47(11-12): 939–943, Dec. 2003.

[11] G. F. Smoot, C. L. Bennett, A. Kogut, et al. Structure in the COBE Differential Microwave Radiometer First-Year Maps. ApJ, 396:L1, Sept. 1992.

[12] D. N. Spergel, L. Verde, H. V. Peiris, et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. ApJS, 148(1):175–194, Sept. 2003.

[13] Planck Collaboration et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. A&A, 641:A1, Sept. 2020.



[14] C. L. Bennett, R. S. Hill, G. Hinshaw, et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission. ApJS, 148 (1):97–117, Sept. 2003.

[15] D. J. Fixsen. The Temperature of the Cosmic Microwave Background. ApJ, 707(2):916–920, Dec. 2009.

[16] Planck Collaboration, N. Aghanim, Y. Akrami, et al. Planck 2018 results. V. CMB power spectra and likelihoods. A&A, 641:A5, Sept. 2020.

[17] B. Wang and J.-J. Wei. An 8.0% Determination of the Baryon Fraction in the Intergalactic Medium from Localized Fast Radio Bursts. ApJ, 944(1): 50, Feb. 2023.

[18] J. P. Gardner, J. C. Mather, M. Clampin, et al. The James Webb Space Telescope. Space Sci. Rev., 123(4):485–606, Apr. 2006.

[19] W. H. Press and P. Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. ApJ, 187:425–438, Feb. 1974.

[20] H. E. Suess and H. C. Urey. Abundances of the Elements. Reviews of Modern Physics, 28(1):53–74, Jan. 1956.

[21] J. H. Jeans. The Stability of a Spherical Nebula. Philosophical Transactions of the Royal Society of London Series A, 199:1–53, Jan. 1902.

[22] S. R. Furlanetto, S. P. Oh, and F. H. Briggs. Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe. Phys. Rep., 433(4-6): 181–301, Oct. 2006.

[23] H. I. Ewen and E. M. Purcell. Observation of a Line in the Galactic Radio Spectrum: Radiation from Galactic Hydrogen at 1,420 Mc./sec. Nature, 168 (4270):356, Sept. 1951.

[24] T. Venumadhav, L. Dai, A. Kaurov, and M. Zaldarriaga. Heating of the intergalactic medium by the cosmic microwave background during cosmic dawn. Phys. Rev. D, 98:103513, Nov 2018.

[25] A. Loeb and M. Zaldarriaga. Measuring the Small-Scale Power Spectrum of Cosmic Density Fluctuations through 21 cm Tomography Prior to the Epoch of Structure Formation. Phys. Rev. Lett., 92:211301, May 2004.

[26] C. M. Hirata. Wouthuysen-Field coupling strength and application to high-redshift 21-cm radiation. Monthly Notices of the Royal Astronomical Society, 367(1):259–274, 03 2006. ISSN 0035-8711.

[27] G. B. Field. Excitation of the Hydrogen 21-CM Line. Proceedings of the IRE, 46:240–250, Jan. 1958.

[28] S. A. Wouthuysen. On the excitation mechanism of the 21-cm (radio-frequency) interstellar hydrogen emission line. Astronomical Journal, 57: 31–32, Jan. 1952.

[29] J. R. Pritchard and S. R. Furlanetto. Descending from on high: Lyman-series cascades and spin-kinetic temperature coupling in the 21-cm line. Monthly

Notices of the Royal Astronomical Society, 367(3):1057–1066, 04 2006.

[30] S. R. Furlanetto. The 21-cm Line as a Probe of Reionization, volume 423, page 247–280. Astrophysics and Space Science Library, 2016.

[31] B. Zygelman. Hyperfine Level-changing Collisions of Hydrogen Atoms and Tomography of the Dark Age Universe. ApJ, 622(2):1356–1362, Apr. 2005.

[32] C. M. Hirata. Wouthuysen-Field coupling strength and application to high-redshift 21-cm radiation. MNRAS, 367(1):259–274, Mar. 2006.

[33] R. Barkana. The rise of the first stars: Supersonic streaming, radiative feedback, and 21-cm cosmology. Phys. Rep., 645:1–59, July 2016.

[34] A. Fialkov, R. Barkana, D. Tseliakhovich, and C. M. Hirata. Impact of the relative motion between the dark matter and baryons on the first stars: semi-analytical modelling. MNRAS, 424(2):1335–1345, Aug. 2012.

[35] J. B. Muñoz. An effective model for the cosmic-dawn 21-cm signal. MNRAS, 523(2):2587–2607, Aug. 2023.

[36] H. T. J. Bevins, W. J. Handley, A. Fialkov, et al. GLOBALEMU: a novel and robust approach for emulating the sky-averaged 21-cm signal from the cosmic dawn and epoch of reionization. MNRAS, 508(2):2923–2936, Dec. 2021.

[37] J. Chardin, G. Kulkarni, and M. G. Haehnelt. Self-shielding of hydrogen in the IGM during the epoch of reionization. MNRAS, 478(1):1065–1076, July 2018.

[38] P. A. Shaver, R. A. Windhorst, P. Madau, and A. G. de Bruyn. Can the reionization epoch be detected as a global signature in the cosmic background? A&A, 345:380–390, May 1999.

[39] T. C. Voytek, A. Natarajan, J. M. Jáuregui García, et al. Probing the Dark Ages at z ~20: The SCI-HI 21 cm All-sky Spectrum Experiment. ApJ, 782 (1):L9, Feb. 2014.

[40] M. e. a. Sokolowski. BIGHORNS - Broadband Instrument for Global HydrOgen ReioNisation Signal. Publications of the Astronomical Society of the Pacific, 32:e004, Feb. 2015.

[41] J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, et al. An absorption profile centred at 78 megahertz in the sky-averaged spectrum. Nature, 555(7694): 67–70, Mar. 2018.

[42] S. Singh, R. Subrahmanyan, N. U. Shankar, et al. SARAS 2: a spectral radiometer for probing cosmic dawn and the epoch of reionization through detection of the global 21-cm signal. Experimental Astronomy, 45(2):269–314, Apr. 2018.

[43] D. C. Price et al. Design and characterization of the Large-aperture Experiment to Detect the Dark Age (LEDA) radiometer systems. MNRAS, 478 (3):4193–4213, Aug. 2018.

[44] S. Singh, N. T. Jishnu, R. Subrahmanyan, et al. On the detection of a cosmic dawn signal in the radio background. Nature Astronomy, 6:607–617, Feb. 2022.

[45] E. de Lera Acedo et al. The REACH radiometer for detecting the 21-cm hydrogen signal from redshift z ≈ 7.5-28. Nature Astronomy, 6:984–998, July 2022.

[46] R. Cappallo, A. Rogers, C. Lonsdale, and J. Bowman. EDGES-3: Deployments and Results. In American Astronomical Society Meeting Abstracts, volume 242 of American Astronomical Society Meeting Abstracts, page 315.04, June 2023.

[47] L. Philip, Z. Abdurashidova, H. C. Chiang, et al. Probing Radio Intensity at High-Z from Marion: 2017 Instrument. Journal of Astronomical Instrumentation, 8(2):1950004, Jan. 2019.

[48] R. A. Monsalve, C. Altamirano, V. Bidula, et al. Mapper of the IGM spin temperature: instrument overview. MNRAS, 530(4):4125–4147, June 2024.

[49] P. Bull, A. El-Makadema, H. Garsden, et al. RHINO: A large horn antenna for detecting the 21cm global signal. arXiv e-prints, arXiv:2410.00076, Sept. 2024.

[50] J. O. Burns, J. Lazio, S. Bale, et al. Probing the first stars and black holes in the early Universe with the Dark Ages Radio Explorer (DARE). Advances in Space Research, 49(3):433–450, Feb. 2012.

[51] G. Hallinan, J. Burns, J. Lux, et al. FARSIDE: A Low Radio Frequency Interferometric Array on the Lunar Farside. In Bulletin of the American Astronomical Society, volume 53, page 379, May 2021. doi: 10.3847/25c2cfeb.60683360.

[52] Y. Shi, F. Deng, Y. Xu, et al. Lunar Orbit Measurement of the Cosmic Dawn's 21 cm Global Spectrum. ApJ, 929(1):32, Apr. 2022.

[53] M. Sathyanarayana Rao, S. Singh, S. K. S., et al. PRATUSH experiment concept and design overview. Experimental Astronomy, 56(2-3):741–778, Sept. 2023.

[54] S. D. Bale, N. Bassett, J. O. Burns, et al. LuSEE 'Night': The Lunar Surface Electromagnetics Experiment. arXiv e-prints, arXiv:2301.10345, Jan. 2023.

[55] R. S. Polidan, J. O. Burns, A. Ignatiev, et al. FarView: An in-situ manufactured lunar far side radio array concept for 21-cm Dark Ages cosmology. Advances in Space Research, 74(1):528–546, July 2024.

[56] X. Chen, F. Gao, F. Wu, et al. Large-scale array for radio astronomy on the farside (LARAF). Philosophical Transactions of the Royal Society of London Series A, 382(2271):20230094, May 2024.

[57] K. Artuc and E. de Lera Acedo. The spectrometer development of CosmoCube, lunar orbiting satellite to detect 21-cm hydrogen signal from

cosmic dark ages. RAS Techniques and Instruments, 4:rzae061, Jan. 2025.

[58] C. D. Brinkerink, M. J. Arts, M. J. Bentum, et al. The Dark Ages Explorer (DEX): a filled-aperture ultra-long wavelength radio interferometer on the lunar far side. arXiv e-prints, arXiv:2504.03418, Apr. 2025.

[59] S. W. Ellingson, T. E. Clarke, A. Cohen, et al. The Long Wavelength Array. IEEE Proceedings, 97(8):1421–1430, Aug. 2009.

[60] C. J. Lonsdale et al. The Murchison Widefield Array: Design Overview. IEEE Proceedings, 97(8):1497–1506, Aug. 2009.

[61] A. R. Parsons et al. The Precision Array for Probing the Epoch of Reionization: Eight Station Results. AJ, 139(4):1468–1480, Apr. 2010.

[62] T.-C. Chang and GMRT-EoR Team. GMRT-EoR Project Update. In American Astronomical Society Meeting Abstracts #221, volume 221 of American Astronomical Society Meeting Abstracts, page 108.04, Jan. 2013.

[63] M. P. van Haarlem et al. LOFAR: The LOw-Frequency ARray. A&A, 556: A2, Aug. 2013.

[64] D. R. DeBoer et al. Hydrogen Epoch of Reionization Array (HERA). PASP, 129(974):045001, Apr. 2017.

[65] B. K. Gehlot et al. The AARTFAAC Cosmic Explorer: observations of the 21-cm power spectrum in the EDGES absorption trough. MNRAS, 499(3): 4158–4173, Dec. 2020.

[66] F. G. Mertens, B. Semelin, and L. V. E. Koopmans. Exploring the Cosmic Dawn with NenuFAR. In A. Siebert, K. Baillié, E. Lagadec, et al., editors, SF2A-2021: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, pages 211–214, Dec. 2021. doi: 10.48550/arXiv.2109.10055.

[67] F. G. Mertens, M. Mevius, L. V. E. Koopmans, et al. Deeper multi-redshift upper limits on the epoch of reionisation 21 cm signal power spectrum from LOFAR between z = 8.3 and z = 10.1. A&A, 698:A186, June 2025.

[68] C. D. Nunhokee, D. Null, C. M. Trott, et al. Limits on the 21 cm Power Spectrum at z = 6.5–7.0 from Murchison Widefield Array Observations. ApJ, 989(1):57, Aug. 2025.

[69] S. Munshi, F. G. Mertens, J. K. Chege, et al. Improved upper limits on the 21-cm signal power spectrum at z = 17.0 and z = 20.3 from an optimal field observed with NenuFAR. MNRAS, 542(4):2785–2807, Oct. 2025.

[70] HERA Collaboration et al. Improved Constraints on the 21 cm EoR Power Spectrum and the X-Ray Heating of the IGM with HERA Phase I Observations. ApJ, 945(2):124, Mar. 2023.

[71] R. Barkana. Possible interaction between baryons and dark-matter particles revealed by the first stars. Nature, 555(7694):71–74, Mar. 2018.

[72] J. B. Muñoz, C. Dvorkin, and A. Loeb. 21-cm Fluctuations from Charged

Dark Matter. Phys. Rev. Lett., 121(12):121301, Sept. 2018.

[73] A. Fialkov and R. Barkana. Signature of excess radio background in the 21-cm global signal and power spectrum. MNRAS, 486(2):1763–1773, June 2019.

[74] I. Reis, A. Fialkov, and R. Barkana. High-redshift radio galaxies: a potential new source of 21-cm fluctuations. MNRAS, 499(4):5993–6008, Dec. 2020.

[75] R. Hills, G. Kulkarni, P. D. Meerburg, and E. Puchwein. Concerns about modelling of the EDGES data. Nature, 564(7736):E32–E34, Dec. 2018.

[76] S. Singh and R. Subrahmanyan. The Redshifted 21 cm Signal in the EDGES Low-band Spectrum. ApJ, 880(1):26, July 2019.

[77] P. H. Sims and J. C. Pober. Testing for calibration systematics in the EDGES low-band data using Bayesian model selection. Monthly Notices of the Royal Astronomical Society, 492(1):22–38, Feb. 2020.

[78] H. T. J. Bevins, W. J. Handley, A. Fialkov, et al. maxsmooth: Rapid Maximally Smooth Function Fitting With Applications in Global 21-cm Cosmology. arXiv e-prints, arXiv:2007.14970, July 2020.

[79] E. Chapman, S. Zaroubi, F. B. Abdalla, et al. The effect of foreground mitigation strategy on EoR window recovery. MNRAS, 458(3):2928–2939, May 2016.

[80] K. Tauscher, D. Rapetti, and J. O. Burns. Formulating and Critically Examining the Assumptions of Global 21 cm Signal Analyses: How to Avoid the False Troughs That Can Appear in Single-spectrum Fits. ApJ, 897(2): 132, July 2020.

[81] D. Anstey, E. de Lera Acedo, and W. Handley. A general Bayesian framework for foreground modelling and chromaticity correction for global 21 cm experiments. Monthly Notices of the Royal Astronomical Society, 506 (2):2041–2058, 06 2021. ISSN 0035-8711. URL https://doi.org/10.1093/mnras/stab1765.

[82] E. Appleton. Wireless studies of the ionosphere. Journal of the Institution of Electrical Engineers, 71:642–650, 1932. URL https://digital-library.theiet.org/doi/abs/10.1049/jiee-1.1932.0144.

[83] R. W. Schunk and A. F. Nagy. Ionospheres: Physics, Plasma Physics, and Chemistry. Cambridge Atmospheric and Space Science Series. Cambridge University Press, 2000.

[84] G. B. Rybicki and A. P. Lightman. Radiative Processes in Astrophysics. Wiley-VCH, 1986.

[85] J. M. Ford and K. D. Buch. Rfi mitigation techniques in radio astronomy. In 2014 IEEE Geoscience and Remote Sensing Symposium, pages 231–234, 2014. doi: 10.1109/IGARSS.2014.6946399.

[86] W. A. Baan. Implementing rfi mitigation in radio science. Journal of Astronomical Instrumentation, 08(01):1940010, 2019.

[87] A. Witze. 'Unsustainable': how satellite swarms pose a rising threat to astronomy. Nature, 606(7913):236–237, May 2022.

[88] D. Anstey and S. A. K. Leeney. Enhanced Bayesian RFI mitigation and transient flagging using likelihood reweighting. RAS Techniques and Instruments, 3(1):372–384, Jan. 2024.

[89] S. A. K. Leeney, W. J. Handley, and E. d. L. Acedo. Bayesian approach to radio frequency interference mitigation. Phys. Rev. D, 108:062006, Sep 2023.

[90] J. H. N. Pattison, D. J. Anstey, and E. de Lera Acedo. Modelling a hot horizon in global 21-cm experimental foregrounds. MNRAS, 527(2):2413–2425, Jan. 2024.

[91] R. F. Bradley, K. Tauscher, D. Rapetti, and J. O. Burns. A Ground Plane Artifact that Induces an Absorption Profile in Averaged Spectra from Global 21 cm Measurements, with Possible Application to EDGES. ApJ, 874(2): 153, Apr. 2019.

[92] O. S. D. O'Hara, Q. Gueuning, E. de Lera Acedo, et al. Uncovering the effects of array mutual coupling in 21-cm experiments with the SKA-Low radio telescope. MNRAS, 538(1):31–48, Mar. 2025.

[93] J. Cumner, C. Pieterse, D. de Villiers, and E. de Lera Acedo. The effects of the antenna power pattern uncertainty within a global 21 cm experiment. MNRAS, 531(4):4734–4745, July 2024.

[94] D. S. Sivia and J. Skilling. Data Analysis - A Bayesian Tutorial. Oxford Science Publications. Oxford University Press, 2nd edition, 2006.

[95] J. Skilling. Nested sampling for general Bayesian computation. Bayesian Analysis, 1(4):833 – 859, 2006. URL https://doi.org/10.1214/06-BA127.

[96] W. J. Handley, M. P. Hobson, and A. N. Lasenby. POLYCHORD: next-generation nested sampling. Monthly Notices of the Royal Astronomical Society, 453(4):4384–4398, Nov. 2015.

[97] R. M. Neal. Slice sampling. The Annals of Statistics, 31(3):705 – 767, 2003. URL https://doi.org/10.1214/aos/1056562461.

[98] T. M. Cover and P. E. Hart. Nearest neighbor pattern classification. IEEE Transactions on Information Theory, 13(1):21–27, 1967.

[99] A. E. Guzmán, J. May, H. Alvarez, and K. Maeda. All-sky Galactic radiation at 45 MHz and spectral index between 45 and 408 MHz. A&A, 525:A138, Jan. 2011.

[100] D. Anstey, E. de Lera Acedo, and W. Handley. Use of time dependent data in Bayesian global 21-cm foreground and signal modelling. MNRAS, 520(1):

850–865, Mar. 2023.

[101] E. Shen, D. Anstey, E. de Lera Acedo, et al. Quantifying ionospheric effects on global 21-cm observations. Monthly Notices of the Royal Astronomical Society, 503(1):344–353, 02 2021. ISSN 0035-8711.

[102] H. K. Vedantham, L. V. E. Koopmans, A. G. de Bruyn, et al. Chromatic effects in the 21 cm global signal from the cosmic dawn. Monthly Notices of the Royal Astronomical Society, 437(2):1056–1069, Jan. 2014.

[103] M. G. Kivelson and C. T. Russell. Introduction to Space Physics. Cambridge University Press, 1995. doi: 10.1017/9781139878296.

[104] S. Chapman. The absorption and dissociative or ionizing effect of monochromatic radiation in an atmosphere on a rotating earth. Proceedings of the Physical Society, 43(1):26–45, jan 1931. URL https://doi.org/10.1088%2F0959-5309%2F43%2F1%2F305.

[105] K. Davies. Ionospheric Radio. Electromagnetic Waves. Institution of Engineering and Technology, 1990.

[106] J. V. Evans and T. Hagfors. Radar astronomy. New York McGraw-Hill, 1968.

[107] J. A. Ratcliffe. An Introduction to the Ionosphere and Magnetosphere. American Journal of Physics, 41(5):761–762, 1973. URL https://doi.org/10.1119/1.1987378.

[108] I. P. Shkarofsky. Generalized Appleton-Hartree Equation for Any Degree of Ionization and Application to the Ionosphere. Proceedings of the IRE, 49 (12):1857–1871, 1961. ISSN 2162-6634.

[109] G. Bernardi et al. A 189 MHz, 2400 deg^2 Polarization Survey with the Murchison Widefield Array 32-element Prototype. ApJ, 771(2):105, July 2013.

[110] C. L. Van Eck, M. Haverkorn, M. I. R. Alves, et al. Diffuse polarized emission in the LOFAR Two-meter Sky Survey. A&A, 623:A71, Mar. 2019.

[111] E. Carretti, M. Haverkorn, L. Staveley-Smith, et al. S-band Polarization All-Sky Survey (S-PASS): survey description and maps. MNRAS, 489(2): 2330–2354, Oct. 2019.

[112] A. Chrysostomou, L. Ball, P. J. Diamond, and et al. SKA1 Operations Plan, Feb. 2020.

[113] M. Spinelli, G. Bernardi, and M. G. Santos. Simulations of Galactic polarized synchrotron emission for Epoch of Reionization observations. MNRAS, 479 (1):275–283, Sept. 2018.

[114] M. Spinelli, G. Bernardi, and M. G. Santos. On the contamination of the global 21-cm signal from polarized foregrounds. MNRAS, 489(3):4007–4015, Nov. 2019.

[115] K. Ferrière, J. L. West, and T. R. Jaffe. The correct sense of faraday rotation.

Monthly Notices of the Royal Astronomical Society, 507(4):4968–4982, 06 2021. ISSN 0035-8711.

[116] B. J. Burn. On the depolarization of discrete radio sources by Faraday dispersion. MNRAS, 133:67, Jan. 1966.

[117] M. A. Brentjens and A. G. de Bruyn. Faraday rotation measure synthesis. A&A, 441(3):1217–1228, Oct. 2005.

[118] J. Cumner, E. de Lera Acedo, D. I. L. de Villiers, et al. Radio Antenna Design for Sky-Averaged 21cm Cosmology Experiments: The REACH Case. Journal of Astronomical Instrumentation, 11(1):2250001-2058, Jan. 2022.

[119] M. Haverkorn, P. Katgert, and A. G. de Bruyn. Structure in the polarized Galactic synchrotron emission, in particular "depolarization canals". A&A, 427:549–559, Nov. 2004.

[120] G. Bernardi, A. G. de Bruyn, M. A. Brentjens, et al. Foregrounds for observations of the cosmological 21 cm line. I. First Westerbork measurements of Galactic emission at 150 MHz in a low latitude field. A&A, 500(3):965–979, June 2009.

[121] E. Lenc et al. Low-frequency Observations of Linearly Polarized Structures in the Interstellar Medium near the South Galactic Pole. ApJ, 830(1):38, Oct. 2016.

[122] C. L. Van Eck, B. M. Gaensler, S. Hutschenreuter, et al. RMTable2023 and PolSpectra2023: Standards for Reporting Polarization and Faraday Rotation Measurements of Radio Sources. ApJS, 267(2):28, Aug. 2023.

[123] A. de Oliveira-Costa, M. Tegmark, B. M. Gaensler, et al. A model of diffuse Galactic radio emission from 10 MHz to 100 GHz. MNRAS, 388(1):247–260, July 2008.

[124] K. M. Górski, E. Hivon, A. J. Banday, et al. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. ApJ, 622(2):759–771, Apr. 2005.

[125] M. McQuinn and R. M. O'Leary. The Impact of the Supersonic Baryon-Dark Matter Velocity Difference on the z ~20 21 cm Background. ApJ, 760 (1):3, Nov. 2012.

[126] A. Mesinger, S. Furlanetto, and R. Cen. 21CMFAST: a fast, seminumerical simulation of the high-redshift 21-cm signal. MNRAS, 411(2):955–972, Feb. 2011.

[127] A. Cohen, A. Fialkov, R. Barkana, and R. A. Monsalve. Emulating the global 21-cm signal from Cosmic Dawn and Reionization. MNRAS, 495(4): 4845–4859, July 2020.

[128] J. A. Vázquez, M. Bridges, M. P. Hobson, and A. N. Lasenby. Model selection applied to reconstruction of the Primordial Power Spectrum. J. Cosmology Astropart. Phys., 2012(6):006, June 2012.



[129] M. Millea and F. Bouchet. Cosmic microwave background constraints in light of priors over reionization histories. A&A, 617:A96, Sept. 2018.

[130] S. Heimersheim, N. S. Sartorio, A. Fialkov, and D. R. Lorimer. What It Takes to Measure Reionization with Fast Radio Bursts. ApJ, 933(1):57, July 2022.

[131] S. Heimersheim, L. Rønneberg, H. Linton, et al. Flexknot and gaussian process for 21 cm global signal analysis and foreground separation, 2023.

[132] M. Bridges, F. Feroz, M. P. Hobson, and A. N. Lasenby. Bayesian optimal reconstruction of the primordial power spectrum. MNRAS, 400(2):1075–1084, Dec. 2009.

[133] J. A. Vázquez, M. Bridges, M. P. Hobson, and A. N. Lasenby. Reconstruction of the dark energy equation of state. J. Cosmology Astropart. Phys., 2012(9):020, Sept. 2012.

[134] J. A. Vázquez, M. Bridges, Y.-Z. Ma, and M. P. Hobson. Constraints on the tensor-to-scalar ratio for non-power-law models. J. Cosmology Astropart. Phys., 2013(8):001, Aug. 2013.

[135] G. Aslanyan, L. C. Price, K. N. Abazajian, and R. Easther. The Knotted Sky I: Planck constraints on the primordial power spectrum. J. Cosmology Astropart. Phys., 2014(8):052–052, Aug. 2014.

[136] K. N. Abazajian, G. Aslanyan, R. Easther, and L. C. Price. The Knotted Sky II: does BICEP2 require a nontrivial primordial power spectrum? J. Cosmology Astropart. Phys., 2014(8):053–053, Aug. 2014.

[137] S. Hee, W. J. Handley, M. P. Hobson, and A. N. Lasenby. Bayesian model selection without evidences: application to the dark energy equation-of-state. MNRAS, 455(3):2461–2473, Jan. 2016.

[138] Planck Collaboration et al. Planck 2015 results. XX. Constraints on inflation. A&A, 594:A20, Sept. 2016.

[139] S. Hee, J. A. Vázquez, W. J. Handley, et al. Constraining the dark energy equation of state using Bayes theorem and the Kullback-Leibler divergence. MNRAS, 466(1):369–377, Apr. 2017.

[140] M. Olamaie, M. P. Hobson, F. Feroz, et al. Free-form modelling of galaxy clusters: a Bayesian and data-driven approach. MNRAS, 481(3):3853–3864, Dec. 2018.

[141] Planck Collaboration et al. Planck 2018 results. X. Constraints on inflation. A&A, 641:A10, Sept. 2020.

[142] L. A. Escamilla and J. A. Vazquez. Model selection applied to reconstructions of the Dark Energy. European Physical Journal C, 83(3):251, Mar. 2023.

[143] W. J. Handley, A. N. Lasenby, H. V. Peiris, and M. P. Hobson. Bayesian inflationary reconstructions from Planck 2018 data. Phys. Rev. D, 100(10):

103511, Nov. 2019.

[144] C. J. Hogan and M. J. Rees. Spectral appearance of non-uniform gas at high z. MNRAS, 188:791–798, Sept. 1979.

[145] D. Scott and M. J. Rees. The 21-cm line at high redshift: a diagnostic for the origin of large scale structure. MNRAS, 247:510, Dec. 1990.

[146] R. Mondal and R. Barkana. Prospects for precision cosmology with the 21 cm signal from the dark ages. Nature Astronomy, 7:1025–1030, Sept. 2023.

[147] J. Dyson. The characteristics and design of the conical log-spiral antenna. IEEE Transactions on Antennas and Propagation, 13(4):488–499, July 1965.

[148] P. Sharma. Astrophysical radio background cannot explain the EDGES 21-cm signal: constraints from cooling of non-thermal electrons. MNRAS, 481 (1):L6–L10, Nov. 2018.

[149] Bell, M. R. and Enßlin, T. A. Faraday synthesis - the synergy of aperture and rotation measure synthesis. A&A, 540:A80, 2012.

