Global 21 cm Signal Recovery Under Changing Environmental Conditions

Joe H. N. Pattison[1,2]⋆, Jean Cavillot [3],†, Harry T.J. Bevins[1,2]‡, Dominic J. Anstey[1,2]§, Eloy de Lera Acedo[1,2]¶

[1]*Astrophysics Group, Cavendish Laboratory, J.J. Thomson Avenue, Cambridge, CB3 0HE, UK*
[2]*Kavli Institue for Cosmology, Madingley Road, Cambridge, CB3 0HA, UK*
[3]*Antenna Group, Université Catholique de Louvain, Louvain-la-Neuve, Belgium*

ABSTRACT

The redshifted 21 cm line of cosmic atomic hydrogen is one of the most auspicious tools in deciphering the early Universe. Recovering this signal remains an ongoing problem for cosmologists in the field, with the signal being hidden behind foregrounds approximately five orders of magnitude brighter than itself. A traditional forward modelling data analysis pipeline using Bayesian data analysis and a physically motivated foreground model to find this signal shows great promise in the case of unchanging environmental conditions. However we demonstrate in this paper that in the presence of a soil with changing dielectric properties under the antenna over time, or a changing soil temperature in the far field of our observation these traditional methods struggle.

In this paper we detail a tool using Masked Auto-regressive Flows that improves upon previous physically motivated foreground models when one is trying to recover this signal in the presence of changing environmental conditions. We demonstrate that with these changing parameters our tool consistently recovers the signal with a much greater Bayesian evidence than the traditional data analysis pipeline, decreasing the root mean square error in the recovery of the injected signal by up to 45%.

Key words: methods: data analysis – early Universe – dark ages – reionization – first stars

1 INTRODUCTION

Between the emission of the cosmic microwave background and the emergence of the first galaxies the 21 cm signal from cosmic hydrogen is expected to be one of the most useful tools in understanding the early Universe. The detection of this signal, however, is an ongoing challenge facing cosmologists. While a complete tomographic map of cosmic hydrogen between cosmic dawn and reionisation remains a distant dream to astrophysicists, the detection of the sky-averaged monopole signal is one closer at hand. A description of the depth, width and general shape of this globally averaged signal would go to unlocking information about early galactic structure formation (Yajima & Li 2014), dark matter (Barkana et al. 2018), and Population III stars (Gessey-Jones et al. 2022).

A potential detection of this signal was made by the Experiment to Detect the Global Epoch of reionisation Signal (EDGES, Bowman et al. 2018), though this detection was in tension with the standard cosmological model - showing a very different shape and depth to our standard theoretical models (e.g. Cohen et al. 2017). Since this detection a number of alternate cosmological models have been proposed to resolve the tensions between the standard model of cosmic evolution and the observed signal, including but not limited to weakly interacting dark matter (e.g. Liu et al. 2019), enhanced radio backgrounds (Fialkov & Barkana 2019; Mittal & Kulkarni 2022), and

competing sources of reheating in the early Universe (Gessey-Jones et al. 2023).

As of yet no follow up experiments have been able to confirm this detection; with Singh et al. (2017) rejecting the EGDES detection with a 95% confidence. In no small part this is the result of the overwhelming foregrounds that obscure the signal from cosmic atomic hydrogen. Galactic foregrounds, predominantly from the synchrotron emission of electrons caught in strong magnetic fields, overwhelms the 21 cm signal by approximately five orders of magnitude, meaning a precise and accurate description of these foregrounds is needed to resolve the signal. This foreground modelling is complicated by a number of factors, including antenna chromaticity (Anstey et al. 2020; Cumner et al. 2023), ionospheric effects (e.g. Shen et al. 2021), man-made radio-frequency interference (RFI) (Leeney et al. 2022; Anstey & Leeney 2023), and a physical horizon (Bassett et al. 2021; Pattison et al. 2024).

All of these sources of signal obfuscation have significant time varying properties which must be corrected for: dielectric properties of the soil in the near field affecting the structure of the beam (Spinelli et al. 2022) which change with Terran weather, solar UV and x-ray emission altering the properties of the ionosphere which will vary with solar weather (Liu et al. 2011), transient RFI sources including plane and satellite radio communication which must be identified and removed (Anstey & Leeney 2023), and changing soil temperatures and dielectric properties in the far field overnight or as seasons progress.

This work expands on Anstey et al. (2023) and Pattison et al. (2024) aiming to account for the time varying properties of the horizon and antenna beam given the use of a physically motivated foreground model.

Section 2 describes how we generate the beams used to describe the antenna in our mock data, discusses how a traditional physically motivated Bayesian data analysis pipeline would function, and discusses the implementation of normalising flows into a novel analysis pipeline to allow signal recovery with changing environmental conditions. Section 3 goes through our results, discussing how the two pipelines recover a 21 cm signal with changing soil moisture levels and temperatures over a series of observations. Section 4 outlines our key conclusions and discusses where we believe this work will lead in future.

2 METHODS

This section details the methods used in this paper to analyse mock data generated in a physically motivated Bayesian data analysis pipeline. The beams used to in this paper are made following the methodology in Section 2.1, and the data simulation process using these beams is discussed in Section 2.2. Sections 2.3 and 2.4 discuss Bayesian statistics and the physically motivated Bayesian forward modelling pipeline used for REACH. Sections 2.5 and 2.6 discuss normalising flows and their application to expand upon the Bayesian analysis pipeline.

2.1 Beam Generation

The beams used in this paper have been obtained after solving the Method of Moments (MoM) (Harrington 1993). In this full-wave solver, the antenna is described using a set of elementary basis functions J_i. Using Galerkin testing, we calculate a matrix of impedance Z_{MoM} which represents the interactions between basis functions and testing functions. Here we assume the presence of a semi-infinite soil with a given permittivity below the antenna. Thus, a spectral formulation, provided in (Cavillot et al. 2019), is adopted where the effect of the soil is accounted for in the calculation of the entries of $\mathbf{Z}_{MoM}$ thanks to reflection coefficients.

In calculating the total current over the antenna J, where J is made by summing the basis functions multiplied by their corresponding coefficients, c_i,

$$J = \sum_i c_i J_i, \tag{1}$$

we must invert the impedance matrix and multiply by the excitation vector to yield our coefficients: $\mathbf{c} = \mathbf{Z}_{MoM}^{-1}\mathbf{v}$. In this case, $\mathbf{v}$ is obtained by applying a delta-gap excitation on the antenna. Once the vector of coefficients $\mathbf{c}$ is obtained, the beams can be generated by performing the Fourier Transform of the current distribution and by multiplying it by the Green's function of the layered medium (Cavillot et al. 2020)

In this case we build a beam analogous to the REACH dipole, which we see in as shown in Figure 1. It is modelled as hexagonal bladed dipole antenna; each blade with dimension of 1400 mm by 928 mm on a 20 x 20 m metallic ground plane with serrated edges, 1 m over a semi-infinite layer of soil (Cumner et al. 2022). We show an example of the beam pattern produced in Figure 2.

2.2 Data Simulation

Simulation of the data we analyse in this paper follows (Pattison et al. 2024). We generate a map of the variation of the spectral indices over the sky using the 2008 Global Sky Model (de Oliveira-Costa et al. 2008). We do this by mapping the 408 MHz onto a corresponding


Figure 1. Overhead image of the REACH hexagonal dipole antenna over its ground plane, situated in the Karoo radio reserve in South Africa. Image courtesy of Dr. Saurabh Pegwal.
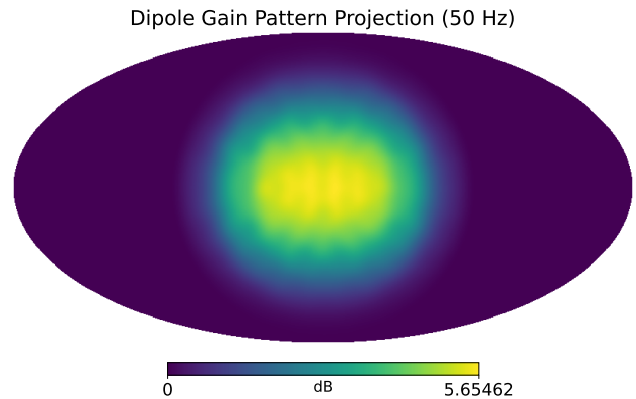

Figure 2. Projection of the antenna gain pattern at 50 MHz of the REACH hexagonal bladed dipole antenna with a ground plane 1 m above a semi infinite layer of soil with a dielectric constant of $\epsilon = 3.8 + 0.29j$.

map at 230 MHz to make sure our map is uncontaminated by the cosmological 21 cm signal.

From this map we subtract any power originating from the cosmic microwave background (CMB), and rotate the map according to the time, date and location of observation. This map is then masked by the horizon surrounding the antenna, generated using the SHAPES algorithm (Bassett et al. 2021), which is given a power equivalent to the expected soil temperature. From here we add additional power to the horizon equal to the power we would expect it to reflect from the sky, and from itself given some reflection coefficient, dependent on its dielectric properties.

Thus our data model reads as

$$T_D(\nu) = \frac{1}{4\pi} \int_0^{4\pi} G(\theta, \phi, \nu) \int_{t_{start}}^{t_{end}} T_{sky}(\theta, \phi, \nu, t) dt d\Omega + \hat{\sigma}, \tag{2}$$

where G is the beam pattern we generate in Section 2.1, θ and ϕ refer to angles over the sky, ν is frequency, t refers to time, $\hat{\sigma}$ is our noise parameter, and the T_{sky} term contains all information regarding the sky maps, CMB and horizon emission and reflection.

For a more comprehensive review of the data generation model used in this paper, see Anstey et al. (2020) and Pattison et al. (2024).

2.3 Bayesian Analysis

Bayesian analysis is a statistical analysis tool derived from Bayes' theorem which aims to improve estimations of probabilities as more data becomes available such that

$$P(\Theta|D, M) = \frac{P(D|\Theta, M)P(\Theta|M)}{P(D|M)} = \frac{\mathcal{L}(\Theta)\pi(\Theta)}{\mathcal{Z}}, \qquad (3)$$

where Θ are the parameters of the model M we use to describe a given dataset D. In the condensed form of this equation the likelihood, $\mathcal{L}(\Theta)$, refers to the probability of observing a dataset, given the estimated parameters and model being used to describe the data, the prior probability, $\pi(\Theta)$, describes our antecedent knowledge of the parameters prior to sampling the data, and the evidence, $\mathcal{Z}$ is the likelihood integrated over all parameters, weighted by the prior.

It is through the calculation of the posterior $P(\Theta|D, M)$ that we are able to estimate the values of the parameters from the prior space. As this space reduces we find our posterior will iteratively become a closer match to the most probable parameter values. The ratio of the evidences of two models will give the ratio of the probability of each model correctly fitting the data, assuming equal prior weights, based on how well the data is described by a given model. To perform this parameter estimation and subsequent model comparison, the REACH pipeline utilises the Nested Sampling algorithm POLY-CHORD (Handley et al. 2015a,b). This algorithm is given a set of prior distributions, from which it will draw a number of samples, and calculate the likelihood of said samples. The samples with the lowest likelihoods are discarded and replaced with a new point drawn from a point of equal or higher likelihood, and the volume of the parameter space from which samples may be drawn from shrinks accordingly. This will continue iteratively until the given termination criterion has been met. This process produces a number of parameter estimations as well as the Bayesian evidence of the model used. For a more in depth discussion of Nested Sampling see Skilling (2006) and Ashton et al. (2022).

2.4 Time-Separated Analysis Pipeline

The global 21-cm signal being 5 orders of magnitude dimmer than the galactic and extragalactic foregrounds means an accurate and precise model of them is required if one wants to extract meaningful information.

This work builds upon a forward modelling method of compensating for foregrounds, originally developed in (Anstey et al. 2020) for use in the Radio Experiment for the Analysis of Cosmic Hydrogen (REACH) (de Lera Acedo et al. 2022). It is a physically motivated model that aims to describe the spatial variation of the foregrounds by fitting for their spectral indices. The sky is divided into a number of regions based on the GSM 230 MHz map (de Oliveira-Costa et al. 2008), such that regions with a similar spectral index are grouped together. One will then fit for the spectral indices of these regions, the temperature and reflection coefficient of the horizon surrounding the antenna (Pattison et al. 2024) and the parameters of 21 cm signal.

With the Earth rotating on its axis through the night and around the Sun through the year the skies overhead are ever changing. These changing skies and the resultant changing foregrounds mean any model that does not account for this time variation will be limited in its ability to extract meaningful information about the signal. Therefore Anstey et al. (2023) further develops this time integrated model to allow for time variation in observation, in which multiple time bins may be fit simultaneously to produce a single set of signal parameters. It is important to note, however, that this model, while being able to account for the movement of the skies overhead assumes that all fitted parameters are constant in each time bin. This is because if we were to let these parameters have time varying properties we


Figure 3. Flowchart showing process for two data bins going through the traditional time-separated analysis pipeline. Inputs are contained in trapeziums, a rectangle denotes a POLYCHORD fit, and data in circles denotes an output. The two bins of data are combined, with a series of temperatures for a given frequency across a range of time bins, these are then fitted with one beam gain pattern (G) to produce a series of nuisance and cosmological parameter estimates and an evidence for the model used.

would see the different parts of the data trying to converge on different posteriors simultaneously.

The likelihood[1] used for this local sidereal time pipeline may be written as:

$$\log\mathcal{L} = \left[\sum_i \sum_j -\frac{1}{2}\log(2\pi\theta_\sigma^2) \right.$$
$$\left. -\frac{1}{2}\left(\frac{T_D(\nu_i, t_j) - (T_F(\nu_i, t_j, \theta_F) + T_S(\nu_i, \theta_S))}{\theta_\sigma} \right)^2 \right], \qquad (4)$$

where $T_D(\nu_i, t_j)$ is the observation data, and $T_F(\nu_i, t_j, \theta_F)$ and $T_S(\nu_i, \theta_S)$ are the foreground and signal models, with θ_σ being a Gaussian noise parameter.

This joint fit across the required number of time bins is performed using POLYCHORD. Throughout this paper this will be henceforth referred to as the 'traditional pipeline'.

Thus for an observation over the course of some period of time the traditional pipeline will combine a number of input data into a matrix spanning over the course of the observation, giving the antenna temperature at each frequency per unit time. Using this data matrix, and a gain pattern of the antenna used to observe this data, we use the traditional pipeline to find a model of the sky, horizon and signal parameters and the corresponding evidences, as shown in Figure 3.

This model must assume that each spectral index, the horizon temperature, and soil reflection coefficient remain constant over the course of an observation. While the assumption that the spectral indices do not vary over time is a robust one, the assumption that the horizon temperature and soil reflection coefficient are unchanging is poor. Thus we must turn to models with greater complexity to allow for this nuance.

2.5 Normalising Flows

To resolve the problem of changing environmental conditions impacting signal recovery we would want to allow these parameters to change during the fitting process. This, however, is not feasible,

---

[1] The specific likelihood used in this work differs from that in Anstey et al. (2023) as the inclusion of a horizon complicates the foreground model. The explicit form and formulation of this likelihood may be found in Appendix A.

allowing all time bins to have parameters independent of one another would lead to an unresolvable tension during the calculation of the posterior. Instead, we try to perform a joint analysis using the parameters that are unchanging across time bins. To do this we aim to perform a fit with samples that we draw from an unknown, joint distribution that we approximate with normalising flows trained on a number of simpler, known, posterior distributions.

Expressive neural networks trained to transform some generic probability distribution into a target distribution are a powerful tool for density estimation. This work utilises the Masked Autoregressive Flows (MAFs) of margarine (Bevins et al. 2022, 2023; Bevins & Handley 2023) to perform this transformation. An MAF is a kind of normalising flow made of a series of Masked Autoencoder for Distribution Estimation (MADE) networks (Germain et al. 2015; Papamakarios et al. 2017).

One of these networks will divide the target distribution into a number of one dimensional conditional probability distributions in which each can be approximated to be a Gaussian.

We draw samples, j, from the complex distribution, θ into our network, outputting values of the mean and standard deviation μ and σ, to reconstruct samples on the standard normal, z, for a given dimension i

$$z_i^j = \frac{\theta_i^j - \mu_i^j}{\sigma_i^j}. \tag{5}$$

As the network goes through its training process the samples on the complex target distribution are transformed onto the standard normal via learnt means and standard deviation. We aim to train the network such that the statistical distance between samples taken from the true distribution and the one produced by the network is minimised. We do this by minimising the Kullback-Leibler divergence between the two distributions, a quantifier of this distance. This is done by trying to maximise the log-probability predicted by the network for a set of samples on the true distribution

$$\underset{w}{\text{argmax}} \sum_{j=0}^{N} \log P_w(\theta^j), \tag{6}$$

where $P_w(\theta)$ is the probability distribution of the network for a given set of weights, w.

Chaining a series of MADEs together we create our MAF, training each link in the chain simultaneously, with each previous iteration feeding into the next. Where given a large enough architecture we can learn complex target distributions.

A more complete description of margarine and the MAFs it uses are found in Bevins et al. (2022, 2023); Bevins & Handley (2023) and Kobyzev et al. (2019); Papamakarios et al. (2017) respectively.

2.6 Novel Joint Analysis Pipeline

The traditional pipeline must assume that soil parameters are unchanging over time, but this is nonphysical. While one may assume that the spectral indices of the sky regions one fits for remain constant, the noise parameter is time independent, and the parameters of the global signal are constant over the course of an observation, the parameters describing the behaviour of the soil will not be. Over the course of a night the soil temperature may vary as much as 30K, and the soil reflection coefficient is heavily dependent on moisture level which is not constant with season. With changing dielectric properties of the soil in the near field we observe significant changes

to the expected beam pattern, which must be accounted for in our analysis of the incoming data.

The novel pipeline combines the traditional pipeline and margarine to allow for the description of these time varying parameters, while maintaining the signal parameters constraining properties of the traditional pipeline in a scenario where the environmental conditions are unchanging. We show this process in Figure 4. Instead of taking all our observation data and combining it for the traditional pipeline to fit with Polychord, we divide it into a series of binned data.

$$T_D(\nu, t) = T_{D_1}(\nu, t) + T_{D_2}(\nu, t) + ... + T_{D_m}(\nu, t), \tag{7}$$

where i and j refer to a How we determine to bin this data will depend on the tolerances of the investigation - an investigation into how overnight soil temperature variation affects signal detection will demand a finer binning than an investigation into antenna variation over series of days.

We divide our data into a series of segments, small enough that the soil properties can be approximated to be constant, but large enough that the pipeline can come near to recovering a signal over a single run. We utilise the traditional pipeline on each of these segments to output a number of parameters Θ_m, their corresponding likelihoods, and an evidence for the overall fit.

These parameters may be further divided in an astrophysical context, we may subdivide our parameter space into nuisance and cosmological parameters of interest, giving $\Theta = \{a, \theta\}$. Here we describe the signal parameters, which in the case of a Gaussian would be depth, width and centre frequency, as our cosmological parameters, and our sky region spectral indices, soil parameters and noise as nuisance parameters. This allows us to marginalise our posterior by integrating out the nuisance parameters, α, such that our new, nuisance-free likelihood reads

$$\mathcal{L}(\theta) = \frac{\int \mathcal{L}(\theta, \alpha)\pi(\theta, \alpha)d\alpha}{\int \pi(\theta, \alpha)d\alpha} = \frac{\mathcal{P}(\theta)\mathcal{Z}}{\pi(\theta)}, \tag{8}$$

as per (Bevins et al. 2022).

We then use margarine to train a series of MAFs on the likelihood distribution from the Polychord outputs from the fits to the individual time bins. Utilising the evidences from the original Polychord fit, we then sample the trained MAFs to give a likelihood weighted mean and standard deviation of each flow. We use these values to produce a new set of priors, such that the priors for each of the parameters are defined as the likelihood weighted mean of the samples ± 5 times the standard deviation. Finally we perform a joint fit with Polychord, combining all these likelihood surfaces into a new likelihood such that

$$\log \mathcal{L}_{1,2,...,m}(\theta) = \log \mathcal{L}_1(\theta) + \log \mathcal{L}_2(\theta) + ... + \log \mathcal{L}_m(\theta). \tag{9}$$

From this we will gain a description of the nuisance-free parameter space. Finally we reduce the evidence the model outputs by the ratio of the volume of the initial set of priors used in the traditional pipeline stage, over the revised priors we input to our second Polychord fit.

3 RESULTS

Here we discuss how the novel and traditional pipelines deal with changing environmental conditions in the near and far field. Section 3.1 describes how the two pipelines are impacted by changing dielectric constants in the nearfield soil, and Section 3.2 details how they deal with changing soil temperatures in the far field.
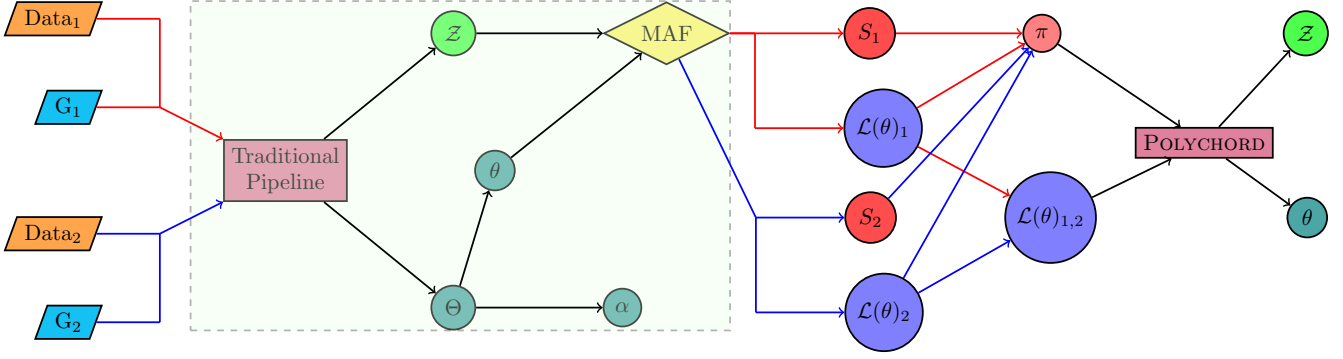
Figure 4. Flowchart showing process for two data bins going through the novel pipeline. A parallelogram represents input data, a rectangle is a Polychord fit, anything in a circle represents an output, and the diamond is the neural network generating normalising flows. The inputs of this pipeline are the antenna temperature data and the gain pattern (G) of the beams used. Each set of inputs are treated separately by everything in the green box, the final output being an MAF, from which we take samples (S) and a likelihood surface. The likelihood surfaces are combined as per Equation 9, and we use the likelihood weighted samples to form a new set of priors for the secondary polychord fit. This fit outputs the revised estimations of the parameters of interest and an evidence for the model.

3.1 Beam Variability

The form of a antenna beam is intrinsically tied to its near field environment - an antenna sitting on a more reflective soil-bed will have a greater directivity upwards, whereas one without this redirected power will look to the horizon more. We explore here how a beam changing over a series of days due to inclement weather will affect the signal recovery using the traditional time separated analysis pipeline with respect to the novel method.

To test this we create a data set using beams where the near-field soil the antenna sits on has three different permittivities, ϵ_i, on three subsequent days. We generate the data as described in Pattison et al. (2024), and inject a Gaussian signal with a central frequency of 85 MHz and a width and depth of 15 MHz and 0.155K respectively, chosen to sit well within the bounds of current theoretical models (Cohen et al. 2017). For this test we use the beam described in Section 2.1, that of the REACH hexagonal dipole antenna 1 m above a bed of single layer, semi-infinite, flat soil, in the Karoo reserve. The permittivites of this soil are such that over the series of three days the dielectric constants are equal to $\epsilon_1 = 7.6 + 0.58j$, $\epsilon_2 = 5.7 + 0.43j$, and $\epsilon_3 = 3.8 + 0.29j$. This kind of scenario would mimic a day of rain over a radiometer followed by two days where the soil gets gradually drier over time. Going forwards we will refer to these as the 'Wet', 'Damp' and 'Dry' beams respectively. An example of the Dry gain pattern is shown in Figure 2. In the signal recovery process we model the sky as having 20 different regions, the spectral indices of which will be fit for by the pipeline to recover the signal parameters.

In Figure 5 we show the extent to which these beam models differ. We plot the average absolute percentage difference between each pixel at each frequency when we compare the Damp and Wet beams to the Dry one. This shows us that in the best case scenario, looking at the frequency channel with the least variation between beams, around 95MHz, we find an average deviation between the Damp and Dry beams, and Damp and Wet beams of approximately 1 and 1.5 parts in 1000. This gets as bad as 3 and almost 6 parts in 1000 at about 52MHz. This means that signal recovery with beams that change to this extent will be challenging when trying to find a signal that sits within the 70-100MHz range that theory predicts (Cohen et al. 2017) and even more difficult with exotic signals.

In testing the traditional pipeline we collate this data into one set, and as this pipeline will only fit our data using a single beam we feed


Figure 5. A comparison of generated beams using different dielectrics to describe the soil under the antenna. We show the mean absolute percentage difference between the Damp and Wet beams at each frequency compared to our Dry beam. The Wet, Damp and Dry beams have dielectric constants of $\epsilon_1 = 7.6 + 0.58j$, $\epsilon_2 = 5.7 + 0.43j$, and $\epsilon_3 = 3.8 + 0.29j$ respectively.

it the Wet, Damp and Dry beams, to test how well it can recover the signal with each of them. We test the novel pipeline by splitting this data into three separate sets according to the beam which produced them, fit them with the correct beams, and then perform a joint fit across the three outputs of the MAFs we train with the pipeline outputs. For this test we assume that we can know precisely the form of the beam. This is a bold assumption, as either it means we can perfectly model the soil, its moisture levels and the extent to which this impacts the gain pattern of the beam, or it means we have a tool that can fit for this beam pattern without fitting out the signal itself. While this approximation may be an optimistic one, we believe it provides a good baseline for an exploration into the power of this novel approach.

We also create a control set, such that the beam is constant over the three days, using the Dry beam. We fit this using the traditional model, comparing RMSEs with the novel pipeline, to show that in

performing a joint fit in the novel method, we maintain the accuracy and precision of the traditional pipeline.

From Table 1 we see that the novel pipeline performs consistently better than the traditional pipeline when we have changing beams. The novel method produces a signal model with far greater Bayesian evidence than any of the traditional models that do not account for our changing beam as the soil dries.

When we compare the novel method to the control, we can see that it performs comparably in terms of RMSE. This shows that the novel method when performing a joint fit over three days with a correctly estimated changing beam performs as well as one might expect a traditional model to given a constant beam over the same time period. We cannot compare the evidences of the control and novel methods as they are fits performed on different data.

Almost all methods are able to recover the signal centre frequency to within twice the quoted error bounds, but to struggle with the signal width and depth. We see that the Wet and Dry beams in Figures 6a and 6c overshoot signal depth. The Wet beam finds the depth incorrectly to 3σ and judges the width incorrectly to the order of 24σ The traditional model using Dry beam finds the Depth to within error, and the width to $\sim 4\sigma$ however, this has consistently the largest error bars of any of the models used. The Damp beam and novel pipeline in Figures 6b and 6d both undershoot the signal depth, and though in terms of absolute magnitude the Damp beam does a worse job of finding the signal depth, the novel pipeline finds this value with a much higher precision, so is further away in terms of the error value quoted.

All models produce a positive $\delta_{\log(\mathcal{Z})}$, is the log of the ratio of the evidences produced when one fits a signal to the data, versus when one chooses to just fit for the foregrounds. This indicates that for all of them, even though recovery is not perfect, we would favour some kind of a detection over a non detection if working with real data. We note here that in the case of the novel pipeline, we describe a fit for just the foregrounds involves fitting for a signal with priors below the level of the noise we inject into our data. We do this as in order to learn the distribution of our signal with MARGARINE we need to provide it with data to learn from. If we solely fit for the foregrounds in the initial stages of the pipeline, there will be no information to work from, so we have no way to compare a set of data with a signal to one without. Restricting the priors to be such that the signal must be below the noise level means that the signal is not recovered, and provides a good surrogate for a foreground only fit.

If one looks solely at the RMSE of the recovered signal compared to the injected one, one might assume that the traditional pipeline recovery that uses the Dry beam in signal recover performs the best. However in practice it has an evidence lower than the traditional method using the Damp beam, and an evidence substantially lower than the novel method. As a result, if this test were to be performed on real data the odds of our novel pipeline giving us the correct signal model would be far higher.

This test is performed not only to demonstrate the capabilities of the novel pipeline, but also to check how the traditional one would deal with recovering a non-exotic 21 cm signal with a changing beam. As the signal becomes more exotic, and drifts further into the regions below 70MHz or above 120MHz we would see the novel pipeline becoming even more powerful as the difference between beams increases further.

3.2 Soil Temperature Variation

Deserts are some of the most extreme environments in the world, where soil temperatures can range between 80°C [2] (Handley 1970) and near freezing. Deserts, such as the Karoo, Inyarrimanha Ilgari Bundara, and Death Valley have been home to the REACH, EDGES and MIST (Monsalve et al. 2024) radiometers, and so we use these as our testing ground for the novel pipeline.

In this section we investigate how the novel and traditional pipelines are able to recover the 21 cm signal in the face of changing soil temperatures, over the course of one evening, and over a year, using the Karoo radio reserve as our test case.

3.2.1 Overnight Temperature Variation

Here we present the results of letting the soil temperature on the horizon vary over the course of 6 hours. We test models with soil changes ranging from an extreme example, to something more realistic, to no variation at all over the course of an evening in order to compare the efficacy of the traditional versus novel pipelines dealing with shifting fitted parameters.

We model the temperature drop overnight as an exponential decay following Newton's law of cooling (Newton 1701), where we assume a constant ambient temperature and a convective heat loss, with

$$T_H = T_a + 40e^{(-kt)}, \tag{10}$$

where T_H is the temperature of the soil on the horizon, T_a is the air temperature after sundown, k a cooling rate constant, and t is the time in hours from midnight. For this test we approximate the air temperature to be 273K after sundown, and vary k between 0 and 0.5 to give examples of slow and rapid temperature reductions over an evening, with k = 0 being a control in which there is no temperature variation. While a more realistic soil temperature model could be found by specific examination of the overnight temperature variation in the Karoo these approximations are used to stress-test the novel pipeline, as opposed to being an exact physical model. This is shown in Figure 7. In the case of the traditional pipeline we observe for the full 5 hours, with a new time bin every 5 minutes; the novel pipeline subdivides this into hour long segments with 5 minute time bins, these outputs are then jointly fit as described in Section 2.6.

From Table 2 and Figure 8 we can see how the traditional and novel methods compare. Both methods find the signal to a high degree of accuracy, with an RMSE corresponding to less than a tenth of the signal depth. However, we can see that both methods show a general trend of increasing RMSE with an increased rate of cooling - the novel method is able to maintain a consistent RMSE for longer, but sees a spike between $k = 0.4$ and $k = 0.5$. Despite both being subject to a higher RMSE at high k values, the novel method recovers a signal consistently closer to the one we inject, with an average RMSE 0.0028 K lower than the traditional one, corresponding to a 20% decrease.

The Bayesian evidence of the signal recovery in the traditional pipeline also notably decreases as k increases with respect to the novel pipeline, which remains mostly consistent with changing k. We see this is Figure 8 where the log of the ratio of the evidences of the traditional pipeline versus the novel one drops dramatically as we increase the cooling constant. The ratio of these evidences indicate

---

[2] There have been reports that natural soil temperatures have reached up to 93.9°C (Kubecka 2001) in Death Valley but these claims are not confirmed.

Table 1. Table detailing signal recovery using the traditional versus novel pipeline when we have a changing antenna beam as the result of decreasing moisture levels in the soil over time. Each model described as 'traditional' follows the signal recovery process described in Section 2.4. The moisture levels beside the model refer to an analogous dielectric constant for the beam we use in signal recovery. The novel recovery model is described in Section 2.6. The inserted mock signal has an 85 MHz Central Frequency, a signal width of 15 MHz, and a Depth of 0.155 K. $\mathcal{Z}_{\mathrm{Gauss}}$ is the Bayesian evidence of trying to fit the injected signal with a Gaussian. $\mathcal{Z}_{\mathrm{No\ 21}}$ is the Bayesian evidence when we try to model for our data having no 21 cm signal. In the traditional model this is done by fitting for only the sky regions and not a signal, in the novel model we restrict the priors in the initial fitting process such that any recovered signal has a depth an order of magnitude below noise level. $\delta_{\mathrm{Log}(\mathcal{Z})}$ is the difference in evidence between these models. RMSE refers to the root mean squared error when comparing the injected mock signal to one that we generate using the posterior averages that our Gaussian model suggests.

	F_0 (MHz)	Signal Width (MHz)	Depth (K)	$\mathrm{Log}(\mathcal{Z}_{\mathrm{Gauss}})$	$\mathrm{Log}(\mathcal{Z}_{\mathrm{No\ 21}})$	$\delta_{\mathrm{Log}(\mathcal{Z})}$	RMSE (K)
Injected Signal	85.0	15.0	0.155				
Models							
Traditional (Wet)	84.1 ± 0.4	10.2 ± 0.2	0.185 ± 0.011	-1468.2 ± 0.4	-1646.6 ± 0.4	178.4 ± 0.6	0.0176
Traditional (Damp)	86.5 ± 0.9	11.1 ± 0.7	0.127 ± 0.015	-477.3 ± 0.4	-555.8 ± 0.4	78.5 ± 0.6	0.0228
Traditional (Dry)	84.3 ± 0.9	12.1 ± 0.7	0.172 ± 0.022	-903.7 ± 0.4	-968.1 ± 0.4	64.4 ± 0.6	0.0101
Novel	85.8 ± 0.5	13.1 ± 0.4	0.133 ± 0.008	7665.7 ± 0.4	7178.3 ± 0.2	487.4 ± 0.4	0.0142
Control	85.9 ± 0.5	13.2 ± 0.4	0.136 ± 0.009	7917.9 ± 0.4	7403.5 ± 0.4	514.4 ± 0.6	0.0130

the odds that one model is correct when compared to the other, so this means that as we increase the cooling constant the odds that the traditional pipeline produces a more correct model than the novel pipeline falls quickly.

3.2.2 Seasonal Temperature Variation

Over the course of a year the average temperature in the Karoo varies by about 10K, at a minimum in July and reaching peaks in January. We investigate how the novel and traditional pipelines are able to deal with this variation by feeding them mock data from observations over the course of a year. We set the temperature over the course of a night to be constant: in January to 290K, in April and October to be 285K and in July to be 280K.

To do this we make four mock data sets, each four hours long with an observation start time as the 1st of each month at midnight. Each of these observations has four one hour time bins. As in the case of the overnight variation, we let the traditional pipeline fit all this data at once, but we let the novel pipeline take each of these days individually to be fit jointly after training the MAFs.

We see from Table 3 that the novel pipeline once again is able to recover the signal to a higher degree of accuracy than the traditional one. The novel pipeline recovers the injected signal with an RMSE 45% lower than that of the traditional pipeline, with a higher Bayesian evidence.

4 CONCLUSIONS

We demonstrate in this paper that a traditional physically motivated Bayesian analysis pipeline like the one designed for REACH will struggle to account for variation in soil moisture levels and the resultant changes to the beam. Thus, the novel pipeline is implemented. This takes the data we get from our antenna and divides it into a number of time bins where we can assume that the soil properties are unchanging. It fits them with the traditional pipeline and trains a series of MAFs from which we take samples to then perform a joint fit across all samples using the POLYCHORD nested sampling algorithm to better recover the redshifted 21 cm signal parameters. We allowing one to fit for signal parameters using changing beams

over time and then using a neural network to collate these data and perform a joint fit across them.

We find that the novel pipeline outperforms the traditional one in the presence of a changing antenna beam. The traditional pipeline can only use one beam model in the fitting process, and with the novel pipeline not having this limitation we see a large improvement. We see that the novel pipeline recovers the injected signal with a log evidence between 7700 and 8800 greater than any of the traditional models, with an average RMSE decrease of 15%. The novel pipeline performs as well with changing beams as we would expect the traditional one to perform with a consistent one, with a comparable RMSE for a given injected signal.

However, the limitation in practice that this pipeline now faces is in working out how soil moisture levels directly impact the dielectrics of a given soil and the subsequent beam. We show in this paper that fitting for our data with a beam that is partially correct over the duration of an observation with a changing beam can increase the RMSE by 75%. Using a beam that is consistently incorrect, whether due to a misunderstanding of dielectric properties, or the topography of the terrain would dramatically decrease likelihood of accurate signal recovery. This can be accounted for by either direct measurement and weather monitoring, or allowing for a parameterized beam fitting. Direct measurement of these soil properties would need to be done on a site by site basis; one needs to take into account moisture draining and retention based on soil types and how soil layers stack on top of one another under the antenna to give a comprehensive understanding of beam evolution over time. A more powerful approach may be to incorporate fitting for the beams gain pattern within the data analysis pipeline. As discussed in Cumner et al. (2023) one can use singular value decomposition to construct a beam pattern using numerical basis functions and fit for these basis functions within the general signal recovery process. While this would add between 30 and 50 dimensions to our likelihood, it could prove to be an enormous improvement in our ability to work through beam uncertainties.

We also discuss how this technique also allows for the pipeline to deal with changing far-field soil temperatures over the course of one night of observation, or changes in average temperature that appear over the course of several months. We show that while the effect of changing farfield soil parameters over these time periods does not have an impact on the ability of the traditional pipeline to recover a signal, we show that it does bias the recovery with
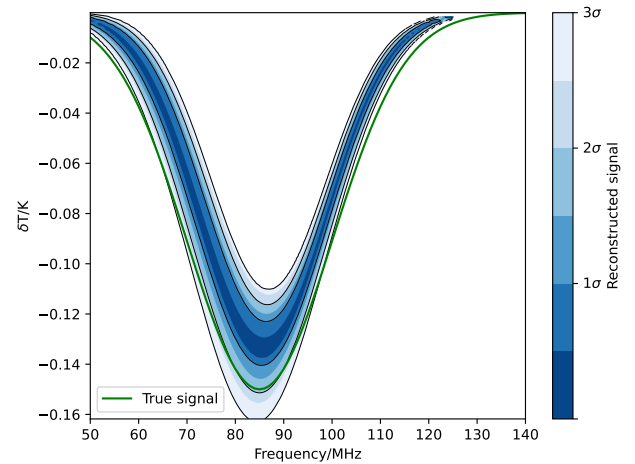
(a) Signal recovery is performed using the traditional pipeline assuming the soil remains in its constant Wet state.
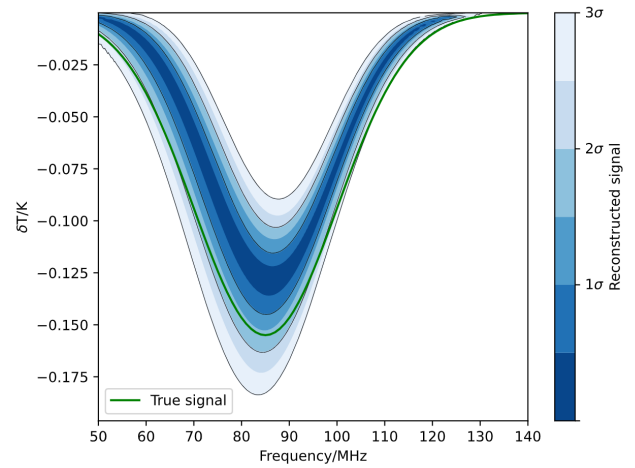
(b) Signal recovery is performed using the traditional pipeline assuming the soil remains in its constant Damp state.

(c) Signal recovery is performed using the traditional pipeline assuming the soil remains in its constant Dry state.

(d) Signal recovery is performed using the novel pipeline where we allow for variation of soil parameters over time.

Figure 6. Recovery of a redshifted 21 cm signal following a 3 day mock observation in which a wet soil gets dryer over time. Injected 'True' signal shown in green, with 85 MHz Central Frequency, 15 MHz Bandwidth, 0.155 K Depth.

respect to the injected signal. The novel pipeline works to recover a signal much closer to that which was injected, showing a 20% RMSE decrease on average over all tests where we let soil temperature change overnight, and an almost 50% RMSE decrease when we tested seasonal temperature variation.

DATA AVAILABILITY

The data that support the findings of this study are available from the first author upon reasonable request.

REFERENCES

Anstey D., Leeney S. A. K., 2023, arXiv e-prints
Anstey D., Acedo E. d. L., Handley W., 2020, Monthly Notices of the Royal Astronomical Society
Anstey D., Acedo E. d. L., Handley W., 2023, Monthly Notices of the Royal Astronomical Society
Ashton G., et al., 2022, Nature Reviews Methods Primers, 2
Barkana R., Outmezguine N. J., Redigolo D., Volansky T., 2018, Physics Review D
Bassett N., Rapetti D., Tauscher K., Nhan B. D., Bordenave D. D., Hibbard J. J., Burns J. O., 2021, The Astrophysical Journal, 923, 33
Bevins H., Handley W., 2023, arXiv e-prints

Table 2. Table comparing signal recovery over a series of mock observations where the temperature of the far field soil varies according to Equation 10 where k describes the cooling constant used in this equation. The novel recovery model is described in Section 2.6. The traditional model follows the signal recovery process described in Section 2.4. The inserted mock signal has an 85 MHz Central Frequency, a signal width of 15 MHz, and a Depth of 0.155 K. $\mathcal{Z}_{\mathrm{Gauss}}$ is the Bayesian evidence of trying to fit the injected signal with a Gaussian. $\mathcal{Z}_{\mathrm{No\ 21}}$ is the Bayesian evidence when we try to model for our data having no 21 cm signal. In the traditional model this is done by fitting for only the sky regions and not a signal, in the novel model we restrict the priors in the initial fitting process such that any recovered signal has a depth an order of magnitude below noise level. $\delta_{\mathrm{Log}(\mathcal{Z})}$ is the difference in evidence between these models. RMSE refers to the root mean squared error when comparing the injected mock signal to one that we generate using the posterior averages that our Gaussian model suggests.

	F_0 (MHz)	Signal Width (MHz)	Depth (K)	$\mathrm{Log}(\mathcal{Z}_{\mathrm{Gauss}})$	$\mathrm{Log}(\mathcal{Z}_{\mathrm{No\ 21}})$	$\delta_{\mathrm{Log}(\mathcal{Z})}$	RMSE (K)
Injected Signal	85.0	15.0	0.155				
$k = 0.0$							
novel	85.4 ± 0.3	13.4 ± 0.3	0.140 ± 0.006	15355.5 ± 0.3	14514.6 ± 0.3	840.9 ± 0.4	0.0107
Traditional	85.7 ± 0.4	13.3 ± 0.3	0.134 ± 0.006	14884.6 ± 0.9	13998.8 ± 0.5	885.8 ± 1.0	0.0135
$k = 0.1$							
Novel	85.5 ± 0.4	13.3 ± 0.3	0.138 ± 0.006	15349.6 ± 0.3	14511.6 ± 0.3	838.0 ± 0.4	0.0115
Traditional	85.8 ± 0.4	13.2 ± 0.3	0.133 ± 0.007	14843.3 ± 0.4	13951.7 ± 0.5	891.6 ± 0.6	0.0142
$k = 0.2$							
Novel	85.5 ± 0.4	13.3 ± 0.3	0.138 ± 0.006	15345.5 ± 0.3	14354.1 ± 0.3	991.4 ± 0.4	0.0115
Traditional	85.9 ± 0.4	13.1 ± 0.3	0.132 ± 0.006	14790.1 ± 0.5	13897.9 ± 0.5	892.2 ± 0.7	0.0145
$k = 0.3$							
Novel	85.5 ± 0.3	13.3 ± 0.3	0.138 ± 0.006	15342.3 ± 0.3	14493.1 ± 0.3	849.2 ± 0.4	0.0119
Traditional	85.9 ± 0.4	13.1 ± 0.3	0.132 ± 0.006	14744.7 ± 0.5	13860.5 ± 0.5	884.2 ± 0.7	0.0147
$k = 0.4$							
Novel	85.6 ± 0.3	13.3 ± 0.3	0.138 ± 0.006	15337.4 ± 0.3	14491.5 ± 0.3	845.9 ± 0.4	0.0118
Traditional	85.9 ± 0.4	13.0 ± 0.3	0.132 ± 0.006	14715.1 ± 0.4	13834.1 ± 0.5	881.0 ± 0.6	0.0150
$k = 0.5$							
Novel	85.6 ± 0.3	13.2 ± 0.3	0.137 ± 0.006	15343.1 ± 0.3	14487.4 ± 0.3	855.7 ± 0.4	0.0123
Traditional	86.0 ± 0.4	13.0 ± 0.3	0.131 ± 0.006	14688.7 ± 0.4	13810.1 ± 0.5	878.6 ± 0.6	0.0151

Table 3. Table comparing signal recovery over a series of mock observations where the temperature of the far field soil varies over the course of a year as described in Section 3.2.2. The novel recovery model is described in Section 2.6. The traditional model follows the signal recovery process described in Section 2.4. The inserted mock signal has an 85 MHz Central Frequency, a signal width of 15 MHz, and a Depth of 0.155 K. $\mathcal{Z}_{\mathrm{Gauss}}$ is the Bayesian evidence of trying to fit the injected signal with a Gaussian. $\mathcal{Z}_{\mathrm{No\ 21}}$ is the Bayesian evidence when we try to model for our data having no 21 cm signal. In the traditional model this is done by fitting for only the sky regions and not a signal, in the novel model we restrict the priors in the initial fitting process such that any recovered signal has a depth an order of magnitude below noise level. $\delta_{\mathrm{Log}(\mathcal{Z})}$ is the difference in evidence between these models. RMSE refers to the root mean squared error when comparing the injected mock signal to one that we generate using the posterior averages that our Gaussian model suggests.

	F_0 (MHz)	Signal Width (MHz)	Depth (K)	$\mathrm{Log}(\mathcal{Z}_{\mathrm{Gauss}})$	$\mathrm{Log}(\mathcal{Z}_{\mathrm{No\ 21}})$	$\delta_{\mathrm{Log}(\mathcal{Z})}$	RMSE (K)
Injected Signal	85.0	15.0	0.155				
Novel	86.6 ± 0.3	12.7 ± 0.4	0.125 ± 0.004	3535.0 ± 0.3	2965.5 ± 0.4	569.6 ± 0.5	0.0187
Traditional	88.0 ± 0.3	10.3 ± 0.2	0.010 ± 0.002	3447.3 ± 1.1	2867.9 ± 0.5	579.4 ± 1.2	0.0337

Bevins H., Handley W., Lemos P., Sims P., de Lera Acedo E., Fialkov A., 2022, in Phys. Sci. Forum.

Bevins H. T. J., Handley W. J., Lemos P., Sims P. H., de Lera Acedo E., Fialkov A., Alsing J., 2023, Monthly Notices of the Royal Astronomical Society, 526, 4613

Bowman J. D., Rogers A. E., Monsalve R. A., Mozdzen T. J., Mahesh N., 2018, Nature, 555, 67

Cavillot J., Tihon D., Craeye C., de Lera Acedo E., Razavi-Ghods N., 2019, in 2019 International Conference on Electromagnetics in Advanced Applications (ICEAA). pp 742–745, doi:10.1109/ICEAA.2019.8879358

Cavillot J., Tihon D., Mesa F., De Lera Acedo E., Craeye C., 2020, IEEE Transactions on Antennas and Propagation, 68, 2753

Cohen A., Fialkov A., Barkana R., Lotem M., 2017, Monthly Notices of the Royal Astronomical Society, 472, 1915

Cumner J., et al., 2022, Journal of Astronomical Instrumentation, 11, 2250001

Cumner J., Pieterse C., De Villiers D., Acedo E. d. L., 2023, arXiv e-prints

Fialkov A., Barkana R., 2019, Monthly Notices of the Royal Astronomical Society

Germain M., Gregor K., Murray I., Larochelle H., 2015, in Bach F., Blei D., eds, Proceedings of Machine Learning Research Vol. 37, Proceedings of the 32nd International Conference on Machine Learning. PMLR, Lille, France, pp 881–889, https://proceedings.mlr.press/v37/germain15.html

Gessey-Jones T., et al., 2022, Monthly Notices of the Royal Astronomical Society, 516, 841

Gessey-Jones T., Fialkov A., Acedo E. d. L., Handley W. J., Barkana R., 2023, arXiv e-prints

Handley N., 1970, Ecology, 51, 434

Handley W. J., Hobson M. P., Lasenby A. N., 2015a, Monthly Notices of the Royal Astronomical Society

Handley W. J., Hobson M. P., Lasenby A. N., 2015b, Monthly Notices of the Royal Astronomical Society, 450, L61

Figure 7. Soil temperature over the course of an evening where each coloured line represents a soil temperature modelled following Equation 10 where the corresponding cooling coefficient for each soil is shown in the legend.

Harrington R. F., 1993, Field Computation by Moment Methods. Wiley-IEEE Press

Kobyzev I., Prince S. J. D., Brubaker M. A., 2019, IEEE Transactions on Pattern Analysis and Machine Intelligence

Kubecka P., 2001, Weather, 56, 218

Leeney S. A. K., Handley W. J., Acedo E. d. L., 2022, arXiv e-prints

Liu L. B., Wan W. X., Chen Y. D., Le H. J., 2011, Solar activity effects of the ionosphere: A brief review, doi:10.1007/s11434-010-4226-9

Liu H., Outmezguine N. J., Redigolo D., Volansky T., 2019, Physical Review D, 100

Mittal S., Kulkarni G., 2022, Monthly Notices of the Royal Astronomical Society, 515, 2901

Monsalve R. A., et al., 2024, Monthly Notices of the Royal Astronomical Society, 530, 4125

Newton I., 1701, Philosophical Transactions of the Royal Society of London, 22

Papamakarios G., Pavlakou T., Murray I., 2017, in Advances in Neural Information Processing Systems. Curran Associates, Inc., doi:10.48550/arXiv.1705.07057

Pattison J. H., Anstey D. J., De Lera Acedo E., 2024, Monthly Notices of the Royal Astronomical Society, 527, 2413

Shen E., Anstey D., De Lera Acedo E., Fialkov A., Handley W., 2021, Monthly Notices of the Royal Astronomical Society, 503, 344

Singh S., Subrahmanyan R., Shankar N. U., Rao M. S., Girish B. S., Raghunathan A., Somashekar R., Srivani K. S., 2017, Experimental Astronomy

Skilling J., 2006, Bayesian Analysis, 1, 833 – 859

Spinelli M., Kyriakou G., Bernardi G., Bolli P., Greenhill L. J., Fialkov A., Garsden H., 2022, Monthly Notices of the Royal Astronomical Society, 515, 1580

Yajima H., Li Y., 2014, Monthly Notices of the Royal Astronomical Society, 445, 3674

de Lera Acedo E., et al., 2022, Nature Astronomy, 6, 984

de Oliveira-Costa A., Tegmark M., Gaensler B. M., Jonas J., Landecker T. L., Reich P., 2008, Monthly Notices of the Royal Astronomical Society, 388, 247

APPENDIX A: LIKELIHOOD FORMULATION

The addition of a horizon to a physically motivated foreground model such that one can fit for soil temperature and reflection coefficient will complicate the likelihood used to recover the signal. In this appendix we describe the formulation and form of the likelihood we use for

signal recovery in a time separated physically motivated Bayesian data analysis pipeline.

As in Equation 4, taken from Anstey et al. (2023), the likelihood takes the generalised form:

$$\log \mathcal{L} = \left[\sum_i \sum_j -\frac{1}{2}\log(2\pi\theta_\sigma^2) \right.$$
$$\left. -\frac{1}{2}\left(\frac{T_D(\nu, t) - (T_F(\nu, t, \theta_F) + T_S(\nu, \theta_S))}{\theta_\sigma} \right)^2 \right], \tag{A1}$$

in which the indices i and j refer to the time and frequency bins. As in Section 2.4, $T_D(\nu, t)$ refers to the observation data, with $T_F(\nu, t, \theta_F)$ and $T_S(\nu, \theta_S)$ are the foreground and signal models respectively, and θ_σ being a Gaussian noise parameter.

The first term in Equation A1 is entirely independent of i and j and thus may be calculated outside the likelihood to give: $-\frac{1}{2}N_t N_\nu \log(2\pi\theta_\sigma^2)$ where N_t and N_ν are the known constants, being number of time and frequency bins respectively.

Putting this to the side for now we focus on the second part of this equation, beginning by restating this section of our likelihood as $\mathcal{L}_{sq}$, and expanding it, allowing us to define

$$L_{sq} = \frac{1}{2\theta_\sigma^2}\left(\left[\sum_i \sum_j T_{D_{i,j}}^2 \right] + \left[\sum_i \sum_j T_{F_{i,j}}(\theta_F)^2 \right] + \right.$$
$$\left[\sum_i \sum_j T_{S_i}(\theta_S)^2 \right] - 2\left[\sum_i \sum_j T_{D_{i,j}}T_{F_{i,j}}(\theta_F) \right]$$
$$\left. -2\left[\sum_i \sum_j T_{D_{i,j}}T_{S_i}(\theta_S) \right] + 2\left[\sum_i \sum_j T_{F_{i,j}}(\theta_F)T_{S_i}(\theta_S) \right] \right). \tag{A2}$$

For ease of parsing we then assign the variables Z, Y, X, W, V, and U to this equation, such that:

$$L_{sq} = \frac{1}{2\theta_\sigma^2}(Z + Y + X - 2W - 2V + 2U), \tag{A3}$$

where throughout this formulation, all final, fully expanded versions of these terms are written in blurple.

We may then expand and evaluate these variables to allow us to maximise the efficiency of the likelihood, calculating what terms we can outside of the likelihood.

We first discuss the terms that do not include $T_{F_{i,j}}$, where our terms will be entirely consistent with (Anstey et al. 2023):

- $Z = \sum_i \sum_j T_{D_{i,j}}^2$, is entirely parameter independent, and is thus a constant, able to be calculated outside of the likelihood.
- $T_{S_i}(\theta_S)$ is time independent, and thus $X = \sum_i \sum_j T_{S_i}(\theta_S)^2$ may be defined as $X = N_t \sum_i T_{S_i}(\theta_S)^2$.
- If one defines $\sum_j T_{D_{i,j}} = T_{D_i}$ then this value may be precalculated, letting $V = \sum_i \sum_j T_{D_{i,j}}T_{S_i}(\theta_S) = \sum_i T_{D_i}T_{S_i}(\theta_S)$.

The remaining terms are all dependent on the foreground model, $T_{F_{i,j}}$, which, when we account for the horizon as in Pattison et al. (2024) becomes:

$$T_{F_{i,j}} = \sum_k K_{i,j,k}F_i(\theta_{F_k}) + \sum_k R_{i,j,k}F_i(\theta_{F_k})|\Gamma|\alpha$$
$$+ J_i T_H (1 + |\Gamma|\alpha). \tag{A4}$$

Figure 8. Plots comparing how the Bayesian evidence and RMSE of signal recovery with the traditional and novel pipeline change over a number of observations with changing soil temperatures. The plot on the left shows the RMSE of the recovered signals for each cooling constant, where the RMSE of the traditional pipeline is shown with a solid line, and the RMSE of the signal recovered with the novel pipeline is shown with a dashed line. The green dotted line indicates an RMSE of one tenth the depth of the injected signal. The Right hand plot shows the log of the ratio of evidences of signal recovery between the traditional and novel pipeline for a dataset with a given cooling constant.

Here k indexes the sky regions in the foreground model. In Anstey et al. (2023) $K_{i,j,k}$ and $F_i(\theta_{F_k})$ are the parameter independent, time dependent term and the parameter dependent time independent term. Following Pattison et al. (2024), with the addition of the horizon around a radiometer $K_{i,j,k}$ specifically becomes a chromaticity term, dealing with power coming from the sky directly; $R_{i,j,k}$ describes power reflected from the horizon. J_i is the analogous horizon 'chromaticity' term, dealing with thermal power originating from the horizon itself, $|\Gamma|$ is the magnitude of the reflection coefficient of the soil, α is defined as $\frac{180-\xi}{180}$ where ξ is the angle the horizon makes relative to the plane of the basin, and T_H is the soil temperature of the horizon.

We then consider the W and U terms:

$$
\begin{aligned}
W &= \sum_i \sum_j T_{D_{i,j}} T_{F_{i,j}}(\theta_F) \\
&= \sum_i \sum_j \sum_k T_{D_{i,j}} K_{i,j,k} F_i(\theta_{F_k}) \\
&+ \sum_i \sum_j \sum_k T_{D_{i,j}} T_{R_{i,j,k}} F_i(\theta_{F_k}) |\Gamma|\alpha \\
&+ \sum_i \sum_j \sum_k T_{D_{i,j}} J_i T_H (1 + |\Gamma|\alpha),
\end{aligned}
\tag{A5}
$$

in which we define $\sum_j T_{D_{i,j}} K_{i,j,k} = T_{D \cdot K_{i,k}}$, and similarly $\sum_j T_{D_{i,j}} R_{i,j,k} = T_{D \cdot R_{i,k}}$; also defining $\sum_j T_{D_{i,j}} = \sum_i T_{D_i}$ means these three terms can be calculated outside the likelihood to speed up computation time to give:

$$
\begin{aligned}
W &= \sum_i \sum_k T_{D \cdot K_{i,k}} F_i(\theta_{F_k}) + \sum_i \sum_k T_{D \cdot R_{i,k}} F_i(\theta_{F_k}) |\Gamma|\alpha \\
&+ \sum_i T_{D_i} J_i T_H (1 + |\Gamma|\alpha).
\end{aligned}
\tag{A6}
$$

When considering the U term:

$$
\begin{aligned}
U &= \sum_i \sum_j T_{S_i}(\theta_s) T_{F_{i,j}}(\theta_F) \\
&= \sum_i \sum_j \sum_k T_{S_i}(\theta_s) K_{i,j,k} F_i(\theta_{F_k}) \\
&+ \sum_i \sum_j \sum_k T_{S_i}(\theta_s) T_{R_{i,j,k}} F_i(\theta_{F_k}) |\Gamma|\alpha \\
&+ \sum_i \sum_j \sum_k T_{S_i}(\theta_s) J_i T_H (1 + |\Gamma|\alpha),
\end{aligned}
\tag{A7}
$$

in which we can see with rearrangement that

$$
\begin{aligned}
\sum_i \sum_j \sum_k T_{S_i}(\theta_s) K_{i,j,k} F_i(\theta_{F_k}) = \\
\sum_i \sum_k \left(\sum_j K_{i,j,k} \right) T_{S_i}(\theta_s) F_i(\theta_{F_k}),
\end{aligned}
\tag{A8}
$$

so if we define $\sum_j K_{i,j,k} = K_{i,k}$, and similarly $\sum_j R_{i,j,k} = R_{i,k}$, both of which can be precalculated we yield the equation for U:

$$
\begin{aligned}
U &= \sum_i \left(\sum_k K_{i,k} F_i(\theta_{F_k}) \right) T_{S_i}(\theta_k) \\
&+ \sum_i \left(\sum_k R_{i,k} F_i(\theta_{F_k}) \right) T_{S_i}(\theta_k) |\Gamma|\alpha \\
&+ N_t \Sigma_i J_i T_H (1 + |\Gamma|\alpha) T_{S_i} F_i(\theta_k)
\end{aligned}
\tag{A9}
$$

Finally, we consider the Y term. To do this we will once again redefine terms from Equation A4, letting $G = \sum_k K_{i,j,k} F_i(\theta_{F_k})$, $H = \sum_k R_{i,j,k} F_i(\theta_{F_k}) |\Gamma|(\alpha)$, and $I = J_i T_H |\Gamma|(1 + \alpha)$, such that:

$$
Y = \sum_i \sum_j T_{F_{i,j}}^2 = G^2 + H^2 + I^2 + 2GH + 2GI + 2HI,
\tag{A10}
$$

in which, analogous to A3, all fully expanded versions of these

terms are written in razzmatazz. We begin our understanding of this equation by expanding G^2:

$$
\begin{aligned}
G^2 &= \sum_i \sum_j \left[\sum_k K_{i,j,k} F_i(\theta_{F_k}) \right]^2 \\
&= \sum_i \sum_k \left[\left[\sum_j K_{i,j,k}^2 \right] F_{i,k}(\theta_{F_k})^2 \right. \\
&\quad \left. + \sum_i \sum_{k_1 \neq k_2} \left[K_{i,j,k_1} K_{i,j,k_1} \right] F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}) \right],
\end{aligned}
\tag{A11}
$$

We will then define $\sum_j K_{i,j,k}^2 = K_{\square i,k}$ and $\sum_j K_{i,j,k_1} K_{i,j,k_2} = K_{\dagger i,k_1,k_2}$, terms which may be calculated outside the likelihood. Thus we may define G^2 as:

$$
\begin{aligned}
G^2 &= \sum_i \sum_k \left[\left[K_{\square i,k} \right]^2 F_{i,k}(\theta_{F_k})^2 \right. \\
&\quad + \sum_{k_1} \sum_{k_2} \sum_i K_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}) \\
&\quad \left. - \mathrm{Tr} \sum_i K_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}} F_{i,K_2}(\theta_{F_{K_2}}). \right.
\end{aligned}
\tag{A12}
$$

As $R_{i,j,k}$ is analogous to $K_{i,j,k}$ we can expand H^2 in the same way such that:

$$
\begin{aligned}
H^2 &= (|\Gamma|\alpha)^2 \sum_i \sum_k \left[\left[R_{\square i,k} \right]^2 F_{i,k}(\theta_{F_k})^2 \right. \\
&\quad + \sum_{k_1} \sum_{k_2} \sum_i R_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}) \\
&\quad \left. - \mathrm{Tr} \sum_i R_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}). \right.
\end{aligned}
\tag{A13}
$$

We continue this process to expand GH, where we allow $\sum_j K_{i,j,k} R_{i,j,k} = RK_{\square i,k}$ and $\sum_j K_{i,j,k_1} R_{i,j,k_2} = RK_{\dagger i,k_1,k_2}$, once again, these terms being able to be calculated outside the likelihood, to give:

$$
\begin{aligned}
GH &= (|\Gamma|\alpha)^2 \sum_i \sum_k \left[\left[RK_{\square i,k} \right]^2 F_{i,k}(\theta_{F_k})^2 \right. \\
&\quad + \sum_{k_1} \sum_{k_2} \sum_i RK_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}) \\
&\quad \left. - \mathrm{Tr} \sum_i RK_{\dagger i,k_1,k_2} F_{i,K_1}(\theta_{F_{K_1}}) F_{i,K_2}(\theta_{F_{K_2}}). \right.
\end{aligned}
\tag{A14}
$$

The final three terms then expand to give:

$$
I^2 = [T_H(1 + |\Gamma|\alpha)]^2 N_t \sum_i J_i^2,
\tag{A15}
$$

$$
GI = T_H(1 + |\Gamma|\alpha) \sum_i \left(\sum_k K_{i,k} J_i F_i(\theta_{F_k}) \right),
\tag{A16}
$$

$$
HI = T_H(|\Gamma|\alpha + (|\Gamma|\alpha)^2) \sum_i \left(\sum_k R_{i,k} J_i F_i(\theta_{F_k}) \right),
\tag{A17}
$$

leaving our final likelihood equation to read as:

$$
\log \mathcal{L} = -\frac{1}{2} N_t N_\nu (2\pi\theta_\sigma^2) - \frac{1}{2\theta_\sigma^2} (Z + Y + X - 2W - 2V + 2U),
\tag{A18}
$$

where all terms are defined above.

This paper has been typeset from a TEX/LATEX file prepared by the author.