



HAL
open science

Implicit Likelihood Inference of Reionization Parameters from the 21 cm Power Spectrum

Xiaosheng Zhao, Yi Mao, Benjamin D. Wandelt

► **To cite this version:**

Xiaosheng Zhao, Yi Mao, Benjamin D. Wandelt. Implicit Likelihood Inference of Reionization Parameters from the 21 cm Power Spectrum. *The Astrophysical Journal*, 2022, 933, <10.3847/1538-4357/ac778e>. <insu-03748246>

HAL Id: insu-03748246

<https://insu.hal.science/insu-03748246v1>

Submitted on 9 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License



Implicit Likelihood Inference of Reionization Parameters from the 21 cm Power Spectrum

Xiaosheng Zhao¹ , Yi Mao¹ , and Benjamin D. Wandelt^{2,3,4} ¹ Department of Astronomy, Tsinghua University, Beijing 100084, People's Republic of China; zhaoxs18@mails.tsinghua.edu.cn, ymao@tsinghua.edu.cn² Sorbonne Université, CNRS, UMR 7095, Institut d'Astrophysique de Paris (IAP), 98 bis bd Arago, F-75014 Paris, France³ Sorbonne Université, Institut Lagrange de Paris (ILP), 98 bis bd Arago, F-75014 Paris, France⁴ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

Received 2022 March 24; revised 2022 May 31; accepted 2022 June 8; published 2022 July 19

Abstract

The first measurements of the 21 cm brightness temperature power spectrum from the epoch of reionization will very likely be achieved in the near future by radio interferometric array experiments such as the Hydrogen Epoch of Reionization Array (HERA) and the Square Kilometre Array (SKA). Standard MCMC analyses use an explicit likelihood approximation to infer the reionization parameters from the 21 cm power spectrum. In this paper, we present a new Bayesian inference of the reionization parameters where the likelihood is implicitly defined through forward simulations using density estimation likelihood-free inference (DELFI). Realistic effects, including thermal noise and foreground avoidance, are also applied to the mock observations from the HERA and SKA. We demonstrate that this method recovers accurate posterior distributions for the reionization parameters, and it outperforms the standard MCMC analysis in terms of the location and size of credible parameter regions. With the minute-level processing time once the network is trained, this technique is a promising approach for the scientific interpretation of future 21 cm power spectrum observation data. Our code `21cmDELFI-PS` is publicly available at this link (<https://github.com/Xiaosheng-Zhao/21cmDELFI>).

Unified Astronomy Thesaurus concepts: [H I line emission \(690\)](#); [Astrostatistics \(1882\)](#); [Bayesian statistics \(1900\)](#); [Neural networks \(1933\)](#); [Reionization \(1383\)](#)

1. Introduction

The cosmic 21 cm background from the epoch of reionization (EoR; Furlanetto et al. 2006) can provide direct constraints on the astrophysical processes regarding how H I gas in the intergalactic medium (IGM) was heated and reionized by the first luminous objects (see, e.g., Dayal & Ferrara 2018; Kannan et al. 2022) that host ionizing sources. Observations of the 21 cm signal with radio interferometric array experiments, including the Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), the Murchison Wide field Array (MWA; Tingay et al. 2013), the LOw Frequency Array (LOFAR; van Haarlem et al. 2013), and the Giant Metrewave Radio Telescope (GMRT; Intema et al. 2017), have focused on the measurements of the power spectrum of the 21 cm brightness temperature fluctuations (hereafter “21 cm power spectrum”) with stringent upper limits on it (Paciga et al. 2013; Pober et al. 2015; Mertens et al. 2020; Trott et al. 2020; Abdurashidova et al. 2022). In the near future, the first measurements of the 21 cm power spectrum from the EoR will very likely be achieved by the Hydrogen Epoch of Reionization Array (HERA; DeBoer et al. 2017) and the Square Kilometre Array (SKA; Mellema et al. 2013) with high signal-to-noise ratio.

The 21 cm power spectrum is a two-point statistic that is sensitive to the parameters in the reionization models (hereafter “reionization parameters”). To shed light on the astrophysical processes during reionization, posterior inference of reionization parameters from future 21 cm power spectrum measurements can

be performed with the Monte Carlo Markov Chain (MCMC) sampling. In the standard MCMC analysis, a multivariate Gaussian likelihood approximation is explicitly assumed, as in the publicly available code `21CMMC` (Greig & Mesinger 2015, 2017, 2018).⁵ Nevertheless, the predefined likelihood approximation may be biased, thereby misestimating the posterior distributions.

To solve this problem, simulation-based inference (SBI; Papamakarios 2019; Cranmer et al. 2020), or so-called “likelihood-free inference” (LFI), is proposed, whereby the likelihood is implicitly defined through forward simulations. This allows building a sophisticated data model without relying on approximate likelihood assumptions. In the Approximate Bayesian Computation (ABC; Cameron & Pettitt 2012; Schafer & Freeman 2012; Hahn et al. 2017), the posterior distribution is approximated by adequate sampling of those parameters that are accepted if the “distance” between the sampling and the observation data meets some criterion. However, the convergence speed of the ABC method is slow in order to get the high-fidelity posterior distribution.

Recently, machine learning has been extensively applied to 21 cm cosmology (e.g., Kern et al. 2017; Shimabukuro & Semelin 2017; Schmit & Pritchard 2018; Gillet et al. 2019; Jennings et al. 2019; Hassan et al. 2020; Choudhury et al. 2022; Zhou & La Plante 2022; Prelogović et al. 2022; Sikder et al. 2022 and references therein). Specifically, Shimabukuro & Semelin (2017) and Doussot et al. (2019) applied neural networks to the estimation of reionization parameters from the 21 cm power spectrum. It is worth noting that such machine-learning applications to 21 cm cosmology are mostly point estimate analyses, i.e., without posterior inference for

⁵ <https://github.com/BradGreig/21CMMC>

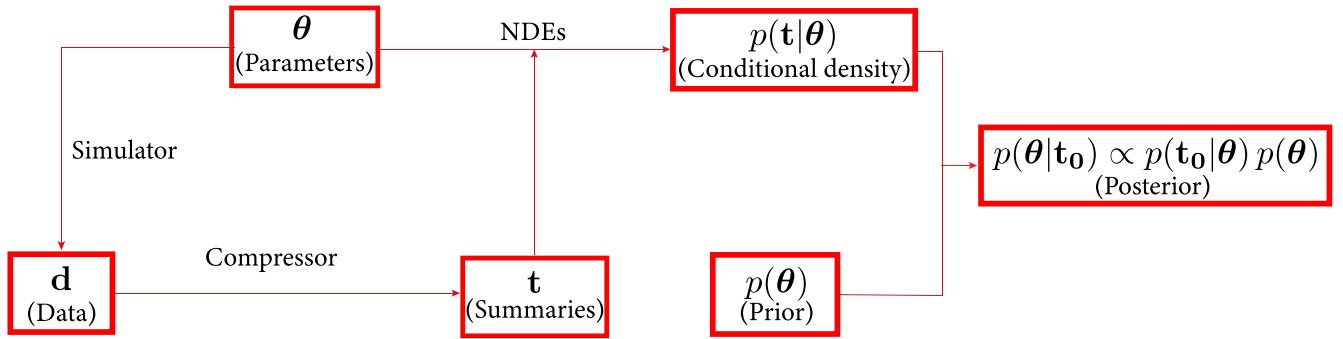


Figure 1. The workflow of 21cmDELFI-PS. Here, the “simulator” refers to 21cmFAST, which generates the 21 cm light-cone data cube (“data”) with the reionization parameters (“parameters”). The “compressor” refers to the procedure of generating the 21 cm power spectra (“summaries”) from the data. The NDEs take the parameter–summary pairs (θ, t) as the input and are trained to learn the conditional density $p(t|\theta)$. The posterior distribution is inferred from the data likelihood evaluated at observation t_0 and parameter prior, using Bayes’ Theorem. Figure adapted from Figure 1 of Z22.

recovered parameters. In Zhao et al. (2022) (hereafter referred to as Z22), we introduced the density estimation likelihood-free inference (DELFI; Alsing et al. 2018, 2019 and references therein) to the 21 cm cosmology, with which the posterior inference of the reionization parameters was performed for the first time from the three-dimensional tomographic 21 cm light-cone images. As a variant of LFI, DELFI contains various neural density estimators (NDEs) to learn the likelihood as the conditional density distribution of the target data given the parameters, from a number of simulated parameter–data pairs. It has been demonstrated to outperform the ABC method in terms of the convergence speed to get the high-fidelity posterior distribution (Alsing et al. 2018).

DELFI is a flexible framework to give the posterior inference of model parameters from data summaries. While the 21 cm power spectra have physical meaning as a summary statistic, from the DELFI point of view, these power spectra are just data summaries of the forward simulations. As such, in this paper, we will apply DELFI in an amortized manner to the problem of posterior inference of reionization parameters from the 21 cm power spectrum. To mock the observations with the HERA and SKA, we will also take into account realistic effects, including thermal noise and foreground avoidance. We will compare the results of DELFI with the standard MCMC analysis using 21cMCMC. To avoid overconfidence in the SBI (Hermans et al. 2021), we will post-validate both marginal and joint posteriors from the SBI (Gneiting et al. 2007; Harrison et al. 2015; Mucesh et al. 2021; Zhao et al. 2021) with statistical tests. To facilitate its application to future observation data, our code, dubbed 21cmDELFI-PS, is made publicly available.⁶

The rest of this paper is organized as follows. We summarize DELFI in Section 2, and describe the data preparation in Section 3, including the simulation of the 21 cm signal and the application of realistic effects. We present the posterior inference results and their validations in Section 4, and make concluding remarks in Section 5. We leave the mathematical definitions of validation statistics to Appendix A, and the effect of the sample size to Appendix B.

2. DELFI Methodology

We summarize DELFI in this section, and refer interested readers to Alsing et al. (2018, 2019) and Z22 for details. DELFI is based on forward simulations that generate data d

given parameters θ . If data vectors are of large dimensions, it is necessary to compress the data d into data summaries t that are of small dimension. DELFI contains various NDEs to learn the conditional density $p(t|\theta)$ from a large number of simulated data pairs $\{\theta, t\}$. NDEs that have been demonstrated to work include mixture density networks (MDN; Bishop 1994) and masked autoregressive flows (MAF; Papamakarios et al. 2017). With the conditional density, the likelihood $p(t_0|\theta)$ can be evaluated at any data summary t_0 from observed data. Then the posterior can be inferred using Bayes’ Theorem, $p(\theta|t_0) \propto p(t_0|\theta)p(\theta)$, where $p(\theta)$ is the prior. The workflow of 21cmDELFI-PS is illustrated in Figure 1.

There are two choices in the application of DELFI: amortized inference and active-learning inference (aka multi-round inference). While it usually needs a large number of preprepared simulations for training the NDEs, amortized inference trains the *global* NDEs only once and the inference from an observation is quick, so post-validations with many mock observations cost only a reasonable amount of computing time. On the other hand, active-learning inference focuses on the most probable region during inference and only trains the NDEs optimized in this local region, so it can effectively save the cost for simulations for inference with data from only *one* observation. However, this “training during inference” process should be repeated for inference with a new set of observational data, so it is computationally expensive to implement post-validations with many mock observations, each using active-learning inference. While the active-learning inference is a valid option in our code package 21cmDELFI-PS, we choose to use the amortized inference in this paper for the purpose of posterior post-validations.

Our code 21cmDELFI-PS employs the PYDELFI⁷ package for the DELFI implementation. We choose the MAFs (see Z22 for a detailed description) as the NDEs with fixed architectures that we find are flexible enough to handle all levels of data sets in 21cmDELFI-PS. For all MAF architectures, we set two neural layers of a single transform, 50 neurons per layer, and five transformations in the MAFs. Ensembles of the MAFs are employed (Alsing et al. 2019; Hermans et al. 2021), and the output posterior is obtained by stacking the posteriors from individual MAFs with weights according to their training errors.

With the trained NDEs, it takes only about 5 minutes with a single core of an Intel Xeon Gold 6248 CPU (base clock speed

⁶ <https://github.com/Xiaosheng-Zhao/21cmDELFI>

⁷ <https://github.com/justinsaling/pydelphi>

2.50 GHz) to process a mock observation and generate the posterior distribution. We plot the posterior contours with the `emcee` module, and run 100 walkers for 1600 steps, with the first 600 steps dropped as “burn-in.”

3. Data Preparation

3.1. Cosmic 21 cm Signal

The 21 cm brightness temperature at position \mathbf{x} relative to the CMB temperature can be written (Furlanetto et al. 2006) as

$$T_{21}(\mathbf{x}, z) = \tilde{T}_{21}(z)x_{\text{HI}}(\mathbf{x})[1 + \delta(\mathbf{x})]\left(1 - \frac{T_{\text{CMB}}}{T_S}\right), \quad (1)$$

where $\tilde{T}_{21}(z) = 27\sqrt{[(1+z)/10](0.15/\Omega_{\text{m}}h^2)(\Omega_{\text{b}}h^2/0.023)}$ in units of mK. Here, $x_{\text{HI}}(\mathbf{x})$ is the neutral fraction, and $\delta(\mathbf{x})$ is the matter overdensity, at position \mathbf{x} . We assume the baryon perturbation traces the cold dark matter on large scales, so $\delta_{\rho_{\text{H}}} = \delta$. In this paper, we focus on the limit where spin temperature $T_S \gg T_{\text{CMB}}$, likely valid soon after reionization, though this assumption is strongly model-dependent. As such, we can neglect the dependence on spin temperature. Also, as a demonstration of concept, we ignore the effect of peculiar velocity; such an effect can be readily incorporated in forward simulations by the algorithm introduced by Mao et al. (2012).

In this paper, we use the publicly available code `21cmFAST`⁸ (Mesinger & Furlanetto 2007; Mesinger et al. 2011), which can be used to perform semi-numerical simulations of reionization, as the simulator to generate the data sets. Our simulations were performed on a cubic box of 100 comoving Mpc on each side, with 66^3 grid cells. Following the interpolation approach in Z22, the snapshots at nine different redshifts of the same simulation box (i.e., with the same initial condition) are interpolated to construct a light-cone 21 cm data cube within the comoving distance of a simulation box along the line of sight (LOS). We concatenate 10 such light-cone boxes, each simulated with different initial conditions in density fields but with the same reionization parameters, together to form a full light-cone datacube of the size $100 \times 100 \times 1000$ comoving Mpc^3 (or $66 \times 66 \times 660$ grid cells) in the redshift range $7.51 \leq z \leq 11.67$. To mimic the observations from radio interferometers, we subtract from the light-cone field the mean of the 2D slice for each 2D slice perpendicular to the LOS, because radio interferometers cannot measure the mode with $k_{\perp} = 0$.

We divide the full light-cone 21 cm datacube into 10 light-cone boxes, each with the size of $(100 \text{ cMpc})^3$ (or 66^3 grid cells), and calculate the *light-cone* 21 cm power spectrum, defined by $\langle \tilde{T}_{21}(\mathbf{k}, z) \tilde{T}_{21}(\mathbf{k}', z) \rangle = (2\pi)^3 \delta(\mathbf{k} + \mathbf{k}') P_{21}(k, z)$. We also use the dimensionless 21 cm power spectrum, $\Delta_{21}^2(k, z) \equiv k^3 P_{21}(k, z) / 2\pi^2$. For each box, we choose to group the modes in Fourier space into 13 k -bins—the upper bound of each k -bin is 1.35 times that of the previous bin. We then combine 10 such power spectra at different central redshifts into a single vector with the size of 130.

We parameterize our reionization model as follows, and refer interested readers to Z22 for a detailed explanation of their physical meanings.

(1) ζ , the ionizing efficiency, which is a combination of several parameters related to ionizing photons. In our paper, we vary ζ as $10 \leq \zeta \leq 250$.

(2) T_{vir} , the minimum virial temperature of halos that host ionizing sources. In our paper, we vary this parameter as $4 \leq \log_{10}(T_{\text{vir}}/\text{K}) \leq 6$.

Cosmological parameters are fixed in this paper as $(\Omega_{\Lambda}, \Omega_m, \Omega_b, n_s, \sigma_8, h) = (0.692, 0.308, 0.0484, 0.968, 0.815, 0.678)$ (Planck Collaboration et al. 2016).

3.2. Thermal Noise

For the thermal noise estimation in this paper, we follow the treatment of the `21CMC` code, for the purpose of comparison on the same ground. The `21CMC` code employs the `21cmsense` module⁹ (Pober et al. 2013, 2014) to simulate the expected thermal noise power spectrum. We summarize the main assumptions here, and refer interested readers to Pober et al. (2013, 2014) and Greig & Mesinger (2015, 2017) for details.

The thermal noise power spectrum of any one mode of uv pixels can be estimated as

$$\Delta_{\text{N}}^2(k) \approx X^2 Y \frac{k^3}{2\pi^2} \frac{\Omega'}{2t} T_{\text{sys}}^2, \quad (2)$$

where $X^2 Y$ is a conversion factor converting observed bandwidths and solid angles to comoving volume in units of $(h^{-1} \text{ Mpc})^3$, Ω' is a beam-dependent factor (Parsons et al. 2014; Pober et al. 2014), and t is the total integration time of all baselines on that particular k -mode. The system temperature $T_{\text{sys}} = T_{\text{rec}} + T_{\text{sky}}$, where T_{rec} is the receiver temperature and T_{sky} is the sky temperature, which can be modeled (Thompson et al. 2001) as $T_{\text{sky}} = 60(\nu/300 \text{ MHz})^{-2.55} \text{ K}$.

The total noise power spectrum for a given k -mode combines the sample variance of the 21 cm power spectrum and the thermal noise using an inverse-weighted summation over all the individual measured modes (Pober et al. 2013; Greig & Mesinger 2017),

$$\delta\Delta_{\text{T+S}}^2(k) = \left(\sum_i \frac{1}{(\Delta_{\text{N},i}^2(k) + \Delta_{21}^2(k))^2} \right)^{-1/2}, \quad (3)$$

where $\delta\Delta_{\text{T+S}}^2(k)$ is the total uncertainty from thermal noise and sample variance in a given k -mode, $\Delta_{\text{N},i}^2(k)$ is the per-mode thermal noise calculated with Equation (2) at each independent k -mode measured by the array as labeled by the index i , and $\Delta_{21}^2(k)$ is the cosmological 21 cm power spectrum (which contributes as the sample variance error here).

In this paper, for `21cmDELFI-PS`, we draw a random noise that follows the normal distribution $N(0, (\delta\Delta_{\text{T+S}}^2(k))^2)$, i.e., with zero mean and the variance of $(\delta\Delta_{\text{T+S}}^2(k))^2$, at each wavenumber k for each sample, and add this noise to the cosmological 21 cm power spectrum in order to generate the mock 21 cm observed power spectrum. As an example, we illustrate the cosmological 21 cm power spectrum, thermal noise, and the mock observation in Figure 2.

In this paper, we consider the mock observations with HERA and SKA, and follow Greig & Mesinger (2015, 2017) for array specifications—except for the minor changes in the

⁸ <https://github.com/andreimesinger/21cmFAST>

⁹ <https://github.com/steven-murray/21cmSense>




Figure 2. An example of the 21 cm power spectrum in the wavenumber range $0.15 \leq k \leq 1.0 \text{ Mpc}^{-1}$ at $z = 7.67$ (left) and $z = 11.47$ (right). Shown are the cosmological light-cone 21 cm power spectrum from the 21cmFAST simulation (red lines) with the reionization parameters defined in Table 2, with the shaded orange regions around it representing the total noise power spectrum (including the contributions from thermal noise and sample variance errors), assuming the measurements with HERA. We also show the thermal noise (black lines), which dominates over the sample variance error, and the mock observed power spectrum (blue dots) with error bars representing the total noise.

SKA design as specified below. For both telescopes, we assume the drift-scanning mode with a total of 1080 hr observation time (Greig & Mesinger 2015). We list the key telescope parameters to model the thermal noise in Table 1.

(1) HERA: we use the core design with 331 dishes forming a hexagonal configuration, with a diameter of 14 m for each dish (Beardsley et al. 2015).

(2) SKA¹⁰: we only focus on the core area with 224 antenna stations randomly distributed within a core radius of about 500 m. The sensitivity improvements due to the arms of the SKA array distribution are negligible.

3.3. Foreground Cut

To remove the bright radio foreground, we adopt the “moderate” foreground avoidance strategy in the 21cmSense module. This strategy (Pober et al. 2014) avoids the foreground “wedge” in the cylindrical $(k_{\perp}, k_{\parallel})$ space, where the “wedge” is defined to extend $0.1 h \text{ Mpc}^{-1}$ beyond the horizon limit (with the slope of about 3.45 at $z = 8$). This also incorporates the coherent addition of all baselines for a given k -mode.

3.4. Database

We use the Latin Hypercube Sampling to scan the reionization parameter space. For the mock observations with only a cosmological 21 cm power spectrum, we generate 18,000 samples with different reionization parameters used for training the NDEs, and 300 additional samples for testing or validating the DELFI. For the mock observations with noise and foreground cut, we generate 9000 samples with different reionization parameters, and make 10 realizations of total noise and foreground cut for each such sample—because the noise is random—so a total of 90,000 samples are used for training the NDEs, and 300 additional samples are used for testing the DELFI. We discuss the effect of the sample size on the

¹⁰ See the SKA1-low configurations in the latest technical document at [this link](#).

Table 1
Specifications for Radio Interferometric Arrays

Parameter	HERA	SKA
Telescope antennas	331	224
Dish diameter (m)	14	35
Collecting area (m ²)	50953	215513
T_{rec} (K)	100	$0.1T_{\text{sky}} + 40$
Bandwidth (MHz)	8	8
Integration time (h)	1080	1080

inference performance in Appendix B. The initial conditions for all samples (with different reionization parameters) were independently generated by sampling spatially correlated Gaussian random fields with the matter power spectrum given by linear theory.

3.5. The 21CMC Setup

For the purpose of comparison, we also run the 21CMC code. We summarize the main setup of 21CMC here, and refer interested readers to Z22 for the details. We generate the mock power spectra at 10 different redshifts, each estimated from a coeval box of 100 comoving Mpc on each side. Strictly speaking, power spectra from the *light-cone* boxes should be employed in both mock observation and MCMC sampling for 21CMC, because the 21cmDELFI-PS is based on the light-cone datacube. However, the 21CMC analysis from the coeval boxes does not significantly change the inference results, since the light-cone effect is only non-negligible at large scales. Since the 21CMC analysis is not the focus of our paper and only serves for comparison, we choose to use coeval boxes in both mock observation and MCMC sampling for 21CMC herein, for self-consistency, to save computational time. This at least avoids the bias caused by using the light-cone box for the mock observation and the coeval box in the MCMC sampling (Greig & Mesinger 2018).

Table 2
Bayesian Inference with 21cmDELFI-PS and 21CMMC for the “Faint Galaxies Model”

Parameter	True value	Pure signal		HERA		SKA	
		21cmDELFI-PS	21CMMC	21cmDELFI-PS	21CMMC	21cmDELFI-PS	21CMMC
$\log_{10}(T_{\text{vir}}/K)$	4.699	$4.699^{+0.011}_{-0.011}$	$4.673^{+0.033}_{-0.025}$	$4.791^{+0.136}_{-0.114}$	$4.779^{+0.128}_{-0.109}$	$4.753^{+0.103}_{-0.095}$	$4.758^{+0.117}_{-0.105}$
$\log_{10}(\zeta)$	1.477	$1.481^{+0.011}_{-0.011}$	$1.444^{+0.026}_{-0.023}$	$1.560^{+0.098}_{-0.076}$	$1.533^{+0.099}_{-0.074}$	$1.540^{+0.074}_{-0.065}$	$1.521^{+0.086}_{-0.073}$

Notes. Here, “Pure signal” refers to the mock observations of cosmological 21 cm power spectrum (i.e., without thermal noise or foreground avoidance); “HERA” (“SKA”) refers to the mock observations of the 21 cm power spectrum from HERA (SKA), which include the total noise (with the contributions from thermal noise and sample variance errors) and foreground avoidance. Note that the fractional error of an observable A is related to the error of its common logarithm by $\Delta A/A \approx 2.3 \Delta \log_{10}(A)$.

Table 3
Same as Table 2, but for the “Bright Galaxies Model”

Parameter	True value	Pure signal		HERA		SKA	
		21cmDELFI-PS	21CMMC	21cmDELFI-PS	21CMMC	21cmDELFI-PS	21CMMC
$\log_{10}(T_{\text{vir}}/K)$	5.477	$5.480^{+0.015}_{-0.016}$	$5.435^{+0.050}_{-0.052}$	$5.364^{+0.097}_{-0.117}$	$5.375^{+0.104}_{-0.144}$	$5.391^{+0.089}_{-0.113}$	$5.379^{+0.102}_{-0.131}$
$\log_{10}(\zeta)$	2.301	$2.306^{+0.023}_{-0.023}$	$2.226^{+0.077}_{-0.072}$	$2.186^{+0.126}_{-0.145}$	$2.159^{+0.133}_{-0.168}$	$2.211^{+0.115}_{-0.141}$	$2.161^{+0.136}_{-0.159}$

In the case with the thermal noise and foreground cut, the total noise power spectrum in Equation (3) is used as the variance of the likelihood function, to include the noise for 21CMMC. In the hypothetical case where there is no thermal noise or foreground in the mock observations, the likelihood function only includes the sample variance from the mock observation, $P_{\text{sv}} = P_{21}(k)/\sqrt{N(k)}$, where $N(k)$ is the number of modes in the spherical shell of k -bin. We turn off the modeling uncertainty parameter in 21CMMC that parameterizes the systematics in the semi-numerical simulations.

We perform the Bayesian inference with 200 walkers for the case of only cosmological 21 cm signal. For each walker, we choose the “burn-in chain” number to be 250, and the main chain number to be 3000. (See the tests in Z22.) For the case with the realistic effects, we employ only 100 walkers with the other settings being the same, because of the convergence of the 21CMMC results.

4. Results

4.1. Posterior Inference for Mock Observations

In this section, we test the Bayesian inference by 21cmDELFI-PS and compare its results with 21CMMC. As a demonstration of concept, we consider two representative mock observations, the “Faint Galaxies Model” and the “Bright Galaxies Model,” whose definitions are listed in the “True value” columns in Tables 2 and 3, respectively, following Greig & Mesinger (2017). These models are chosen as two examples with extreme parameter values. Their global reionization histories are similar, but reionization in the “Faint Galaxies Model” is powered by more abundant low-mass galaxies yet with smaller ionization efficiency (due to the smaller escape fraction of ionizing photons) than in the “Bright Galaxies Model,” so the H II bubbles in the former are smaller and more fractal than in the latter. In Tables 2 and 3, we list the median and 1σ errors of posterior inference for both mock observations, respectively.

In the hypothetical case where there is no thermal noise or foreground in the mock observations of the cosmological 21 cm power spectrum, we show the results of posterior inference for

both mock observations in Figure 3. Both 21cmDELFI-PS and 21CMMC can recover posterior distributions for the reionization parameters in the sense that the medians are within the estimated 1σ confidence region. However, 21cmDELFI-PS outperforms 21CMMC in terms of the location and size of credible parameter regions. Quantitatively, for the “Faint Galaxies Model,” the systematic shift (i.e., relative errors of the predicted medians with respect to the true values) and the 1σ statistical errors are $0\% \pm 0.23\%$ ($-0.55\%^{+0.70\%}_{-0.53\%}$) for $\log_{10}(T_{\text{vir}})$ with 21cmDELFI-PS (21CMMC), respectively, and $0.27\% \pm 0.74\%$ ($-2.2\%^{+1.8\%}_{-1.6\%}$) for $\log_{10} \zeta$ with 21cmDELFI-PS (21CMMC), respectively. Not only are the predicted medians of 21cmDELFI-PS much closer to the true values than 21CMMC, but also the estimated statistical errors in the former are 2–3 times smaller than in the latter. These results hold generically for the “Bright Galaxies Model,” as well. In 21CMMC, an explicit likelihood assumption is made, namely that the likelihood is a multivariate Gaussian, with independent measurements at each redshift and at each k -mode. Our results question the validity of this explicit likelihood assumption in the hypothetical case where there is no thermal noise or foreground in the mock 21 cm data. Indeed, Mondal et al. (2016), Shaw et al. (2019), and Shaw et al. (2020) show that the non-Gaussianity in the covariance of the 21 cm power spectrum is non-negligible.

Now we apply the realistic effects, including the total noise power spectrum (with thermal noise and sample variance errors), in addition to using a foreground cut to avoid the foreground. In Figures 4 and 5, we show the results of posterior inference for mock observations with HERA and SKA, respectively. We find that both 21cmDELFI-PS and 21CMMC can recover posterior distributions for the reionization parameters in this case. Also, the performances with HERA and SKA are comparable, which was also found in Greig & Mesinger (2015, 2017). For mock observations with both HERA and SKA, the statistical (fractional) errors are $\sim 2\%–3\%$ for $\log_{10}(T_{\text{vir}})$ and $\sim 5\%–7\%$ for $\log_{10} \zeta$, or equivalently $\sim 23\%–33\%$ for T_{vir} and $\sim 17\%–39\%$ for ζ .

Comparing 21cmDELFI-PS and 21CMMC, their recovered medians are in good agreement, but the 1σ credible regions

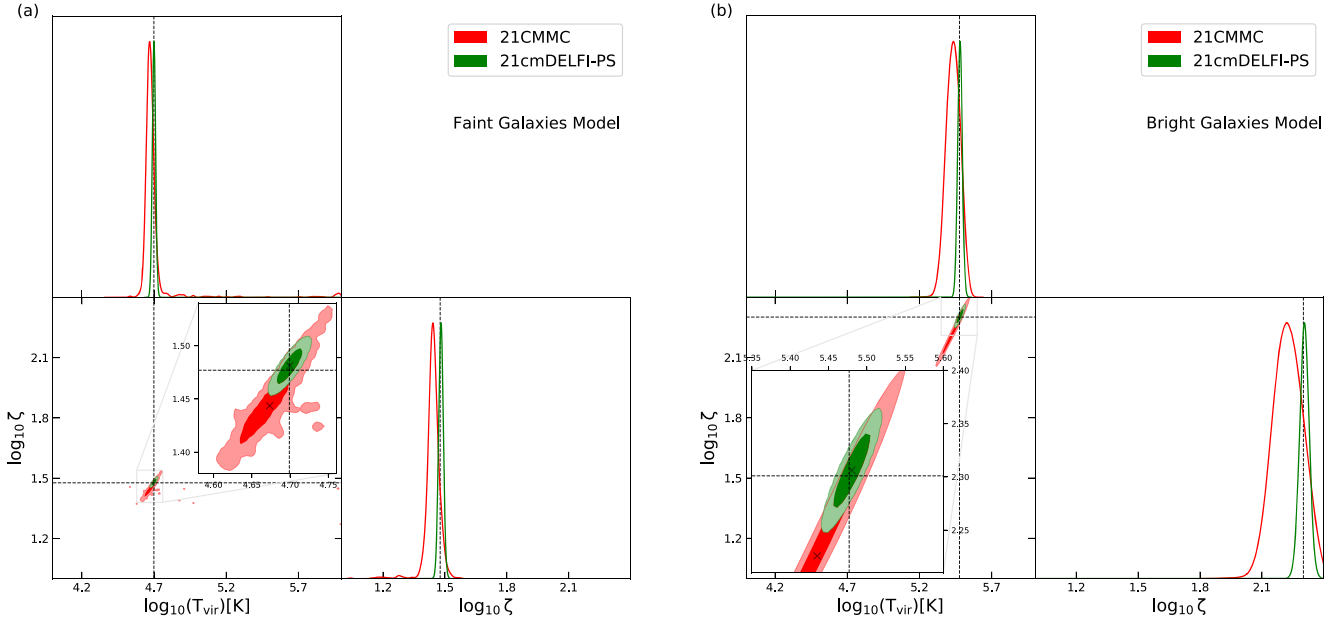


Figure 3. The posteriors estimated from the cosmological 21 cm power spectrum (i.e., without thermal noise or foreground avoidance) by two different approaches, 21cmDELFI-PS (green) and 21CMMC (red), for two mock observations: the “Faint Galaxies Model” (left) and the “Bright Galaxies Model” (right). We show the median (cross) as well as the 1σ (dark) and 2σ (light) confidence regions. The dashed lines indicate the true parameter values.




Figure 4. Same as Figure 3, but the estimations are made from mock observations of the 21 cm power spectrum from HERA, which include the total noise (with the contributions from thermal noise and sample variance errors) and foreground avoidance.

estimated by 21cmDELFI-PS are in general slightly smaller than those by 21CMMC. This agreement suggests that the explicit Gaussian likelihood assumption is approximately valid when thermal noise and foreground cut are incorporated in the mock 21 cm data. This is reasonable because thermal noise dominates over the cosmic variance in HERA and SKA, and thus the non-Gaussian covariance due to the cosmic variance is subdominant. Also, we find that the directions of degeneracies in the parameter space are almost the same for these two codes, which reflects the fact that both codes use the same simulator, 21cmFAST, so the dependency of the 21 cm power spectrum on the reionization parameters is the same.

It is worthwhile to note that the mock observed power spectrum used as the input of 21CMMC is the cosmological 21 cm signal from the 21cmFAST simulation (i.e., the red line in Figure 2), which is the default setting of 21CMMC. If we apply noise to the mock observed power spectrum (i.e., the blue dots in Figure 2) used as the input of 21CMMC, the inference performance of 21CMMC can be degraded significantly. However, this is technically solvable by generating multiple noise realizations in the MCMC sampling, albeit with considerably larger computational cost. In comparison, the input of 21cmDELFI-PS is the mock with noise, in which case the Bayesian inference works well. This is because




Figure 5. Same as Figure 4, but with mock observations of the 21 cm power spectrum from SKA.

21cmDELFI-PS learns the effect of variations due to noise by including 10 noise realizations for each reionization model in the training samples, with reasonable computational time for training.

4.2. Validation of the Posterior

In this subsection, we perform the validation¹¹ of posteriors, which tests statistically the accuracies of inferred posterior distributions. We employ 300 samples for mock observations with only cosmological 21 cm power spectrum and with realistic effects with HERA and with SKA, respectively. These samples are randomly chosen from the allowed region in the parameter space in which the neutral fraction satisfies $0.08 \leq x_{\text{HI}} \leq 0.81$ at $z = 7.1$, corresponding to the 2σ confidence region from the IGM damping wing of ULASJ1120 + 0641 (Greig et al. 2017). The mathematical definitions of validation statistics are left to Appendix A.

(1) Validation of marginal posteriors. We first focus on the validation of posteriors for each single parameter marginalized over other parameters. In Figure 6, we perform two statistics—(i) quantile of the probability integral transform (PIT), and (ii) $\hat{F}_I(\theta) - \tilde{G}_I(\theta)$, where $\hat{F}_I(\theta)$ is the predictive cumulative distribution function (CDF) and $\tilde{G}_I(\theta)$ is the empirical CDF, for $\theta = \log_{10}(T_{\text{vir}})$ and $\log_{10} \zeta$. In the quantile–quantile (QQ) plot, we find that the curves for all cases of mock observations are close to the diagonal line ($Q_{\text{data}} = Q_{\text{theory}}$). In the marginal calibration, we find that the values of $\hat{F}_I(\theta) - \tilde{G}_I(\theta)$ for all cases of mock observations are small (less than 0.1). These findings meet the expectations with which the marginal posterior distributions are probabilistically calibrated.

(2) Joint posteriors validation. Next, we focus on the validation of posteriors in the joint parameter space. In Figure 7, we consider three statistics: (i) quantile of the copula probability integral transformation (copPIT); (ii) $\hat{\mathcal{K}}_{H_I}(w) - \tilde{J}_I(w)$, where $\hat{\mathcal{K}}_{H_I}(w)$ is

the average Kendall distribution function and $\tilde{J}_I(w)$ is the empirical CDF of the joint CDFs, and where w is a variable between zero and unity; and (iii) the highest probability density (HPD). In both QQ plots for the copPIT and for the HPD, we find that the curves for all cases of mock observations are close to the diagonal line ($Q_{\text{data}} = Q_{\text{theory}}$). In the Kendall calibration, we find that the values of $\hat{\mathcal{K}}_{H_I}(w) - \tilde{J}_I(w)$ for all cases of mock observations are small (less than 0.1). These findings meet the expectations with which the joint posterior distributions are probabilistically copula calibrated and probabilistically HPD calibrated.

(3) Hypothesis tests. The aforementioned validations of marginal and joint posteriors provide qualitative measures of the validation of the inferred posteriors. The quantitative metrics for the uniformity of distribution of these statistics (PIT, copPIT, and HPD) are in the hypothesis tests. For both marginal posteriors and joint posteriors, we perform the Kolmogorov–Smirnov (KS) test and the Cramér–von Mises (CvM) test. The p -value of the KS (CvM) test is the probability to obtain a sample data distribution with D (D^*) larger than the measured D_{obs} (D_{obs}^*). The difference between these two metrics is that the KS measure D is sensitive to the median, but the CvM measure D^* incorporates the information of the tails of a distribution. In Table 4, we list the p -values for the PIT of two individual reionization parameters (for the validation of marginal posteriors), and those for the copPIT and the HPD (for the joint posteriors validation). In all cases of mock observations, we find that the p -value for the PIT and for the copPIT is larger than 0.05, and that for the HPD is larger than 0.01. This means that, at the significance level of 5%/5%/1%, the null hypothesis that the PIT/copPIT/HPD distribution is uniform is accepted. In other words, the posteriors are probabilistically calibrated, probabilistically copula calibrated, and probabilistically HPD calibrated.

In sum, the posteriors obtained with 21cmDELFI-PS are valid under all statistical tests of marginal and joint posteriors.

¹¹ In the convention of statistics, these tools are called “calibration.” In our context of testing the posteriors, however, the term “validation” is more appropriate.

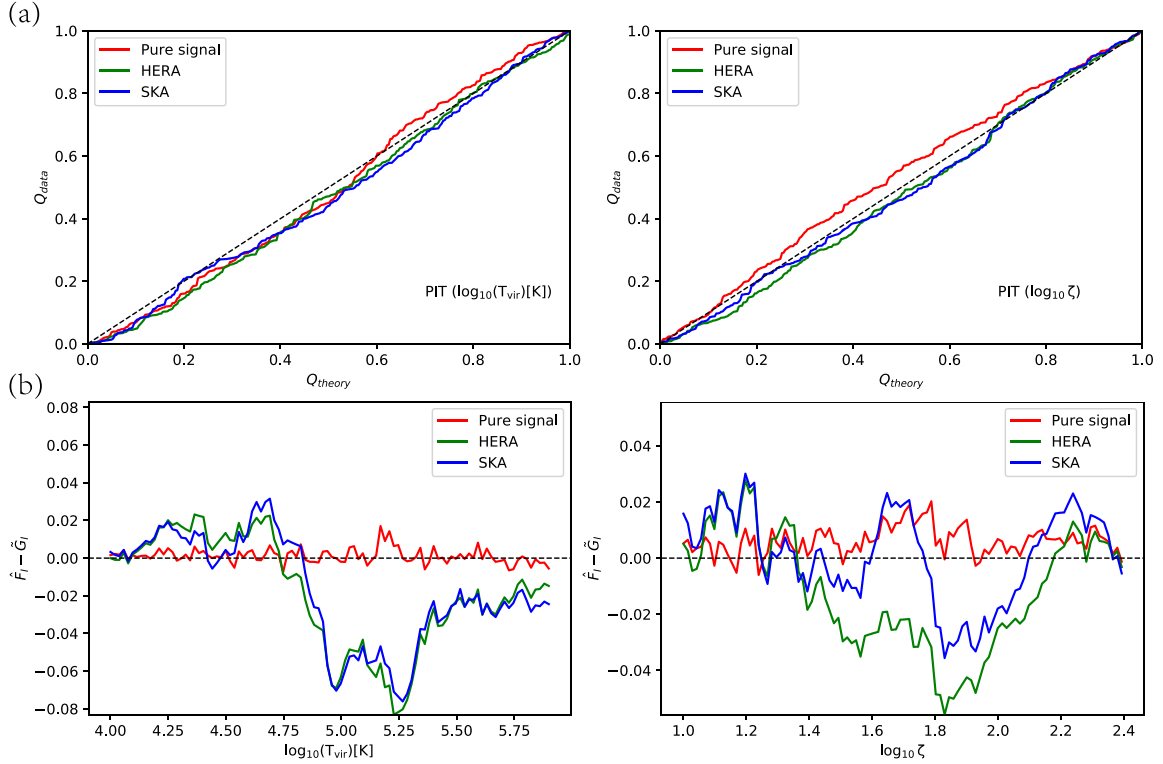


Figure 6. Validation of marginal posteriors. (a) Quantile–quantile plot, i.e., quantile of the PIT distribution from the test sample data vs. that from a theoretical uniform PIT distribution. The diagonal dashed line represents the ideal case ($Q_{\text{data}} = Q_{\text{theory}}$). (b) Marginal calibration, i.e., $\hat{F}_j(\theta) - \hat{G}_j(\theta)$. The diagonal dashed line represents the ideal case (identically zero). We show the results for $\theta = \log_{10}(T_{\text{vir}})$ (left panels) and $\log_{10}\zeta$ (right panels), respectively, and for mock observations with only cosmological 21 cm signal (“Pure signal”; red), HERA (green), and SKA (blue), respectively.




Figure 7. Joint posteriors validation. (a) Quantile–quantile plot, i.e., quantile of the copPIT distribution from the test sample data vs. that from a theoretical uniform copPIT distribution. (b) Kendall calibration, i.e., $\hat{K}_{H_j}(w) - \hat{J}_j(w)$, where w is a variable between zero and unity. The diagonal dashed line represents the ideal case (identically zero). (c) Quantile–quantile plot, but for the quantile of the HPD distribution. We show the results for mock observations with only cosmological 21 cm signal (“Pure signal”; red), HERA (green), and SKA (blue), respectively. In panels (a) and (c), the diagonal dashed line represents the ideal case ($Q_{\text{data}} = Q_{\text{theory}}$).

Table 4

The p -values for the Null Hypotheses That These Statistics Are of a Uniform Distribution

Statistics	Pure signal		HERA		SKA	
	KS	CvM	KS	CvM	KS	CvM
PIT ($\log_{10}(T_{\text{vir}})$)	0.37	0.13	0.16	0.09	0.11	0.10
PIT ($\log_{10}\zeta$)	0.14	0.05	0.30	0.20	0.36	0.36
copPIT	0.54	0.39	0.13	0.17	0.22	0.29
HPD	0.04	0.01	0.03	0.02	0.73	0.67

5. Summary

In this paper, we present a new Bayesian inference of the reionization parameters from the 21 cm power spectrum. Unlike the standard MCMC analysis that uses an explicit

likelihood approximation, in our approach, the likelihood is implicitly defined through forward simulations using DELFI. In DELFI, once the neural density estimators are trained to learn the likelihood, the inference speed for a test observation of power spectrum is very fast—about 5 minutes with a single core of an Intel CPU (base clock speed 2.50 GHz).

We show that this method (dubbed 21cmDELFI-PS) recovers accurate posterior distributions for the reionization parameters, using mock observations with and without realistic effects of thermal noises and foreground cut with HERA and SKA. For the purpose of comparison, we perform MCMC analyses of the 21 cm power spectrum using a Gaussian likelihood approximation. We demonstrate that this new method (21cmDELFI-PS) outperforms the standard MCMC analysis (using 21cMCMC) in terms of the location and size of

credible parameter regions, in all cases of mock observations. We also perform the validation of both marginal and joint posteriors with sophisticated statistical tests, and demonstrate that the posteriors obtained with 21cmDELFI-PS are statistically self-consistent.

It is interesting to note that, in the scenario of only cosmological 21 cm signal without thermal noises and foreground contamination, the 21 cm power spectrum analysis with 21cmDELFI-PS even outperforms our previous analysis from 3D tomographic 21 cm light-cone images with DELFI-3D CNN (Z22), where 3D CNN was applied to compress the light-cone images into low-dimensional summaries. Since the power spectrum only contains a subset of information in images, this implies that 3D CNN is not the optimal imaging compressor. We leave it to a follow-up of Z22 to search for a new imaging compressor with better performance results than 3D CNN and 21cmDELFI-PS.

The development of this new Bayesian inference method is timely because the first measurements of the 21 cm power spectrum from the EoR will very likely be achieved in the near future by HERA and SKA. The DELFI framework is flexible with regard to incorporating more realistic effects through forward simulations, so this technique will be a promising approach for the scientific interpretation of future 21 cm power spectrum observation data. To facilitate its application, we have made 21cmDELFI-PS publicly available.

As a demonstration of concept, this paper only considers the limit where spin temperature $T_S \gg T_{\text{CMB}}$, and therefore neglects the dependence on spin temperature. We leave it to future work to extend the parameter space to the cosmic dawn parameters that affect the IGM heating and Ly α -pumping. Also, Bayesian inference of the reionization parameters from other summary statistics, e.g., 21 cm bispectrum (Watkinson et al. 2022), can be developed in a manner similar to that of 21cmDELFI-PS, because these statistics are just data summaries from the DELFI point of view. We leave it to future work to make such developments.

This work is supported by National SKA Program of China (grant No. 2020SKA0110401), NSFC (grant No. 11821303), and National Key R&D Program of China (grant No. 2018YFA0404502). B.D.W. acknowledges support from the Simons Foundation. We thank Paulo Montero-Camacho, Jianrong Tan, Steven Murray, Nicholas Kern and Biprateep Dey for useful discussions and help, and the anonymous referee for constructive comments. We acknowledge the Tsinghua Astrophysics High-Performance Computing platform at Tsinghua University for providing computational and data storage resources that have contributed to the research results reported within this paper.

Software: 21CMMC (Greig & Mesinger 2015, 2017, 2018), 21cmFAST (Mesinger & Furlanetto 2007; Mesinger et al. 2011), pydelfi (Alsing et al. 2019), TensorFlow (Abadi et al. 2016), GetDist (Lewis 2019), NumPy (Harris et al. 2020), Matplotlib (Hunter 2007), SciPy (Virtanen et al. 2020), scikit-learn (Pedregosa et al. 2011), Python2 (Van Rossum & Drake 1995), Python3 (Van Rossum & Drake 2009), galpro (Mucesh et al. 2021), seaborn (Waskom 2021), Astropy (Astropy Collaboration et al. 2013, 2018).

Appendix A Statistical Tools for Validations

In this section, we introduce some statistical tools for validation of the marginal or joint posterior distributions inferred from data, following Gneiting et al. (2007), Ziegel & Gneiting (2014), Harrison et al. (2015), and Mucesh et al. (2021).

A.1. Validation of Marginal Posteriors

The tools in this subsection focus on the validation of posteriors for each single parameter marginalized over other parameters.

A.1.1. Probabilistic Calibration

Consider an inferred marginal distribution $f(\theta)$, where θ is a marginalized parameter. For example, $\theta = \log_{10}(T_{\text{vir}})$ or $\log_{10} \zeta$ in this paper, and $f(\theta)$ is the probability distribution function that is generated in the MCMC chain given the observed data. We define the probability integral transform (PIT) as the cumulative distribution function (CDF) of this marginal distribution (Gneiting et al. 2007; Mucesh et al. 2021),

$$\text{PIT}(\tilde{\theta}) \equiv \int_{-\infty}^{\tilde{\theta}} f(\theta) d\theta, \quad (\text{A1})$$

where $\tilde{\theta}$ is the true value given the data. The marginal distributions are probabilistically calibrated if true values are randomly drawn from the real distributions. This is equivalent to the statement that the distribution of PIT is uniform. In the so-called quantile-quantile (QQ) plot, wherein the quantile of the PIT distribution from the data is compared with that from a uniform distribution of PIT, this QQ plot falls on the diagonal line if the PIT distribution is exactly uniform.

A.1.2. Marginal Calibration

The uniformity of PIT distribution is only a necessary condition for real marginal posteriors (Hamill 2001; Gneiting et al. 2007; Mucesh et al. 2021; Zhao et al. 2021). As a complementary test, the marginal calibration (Gneiting et al. 2007; Mucesh et al. 2021) compares the average predictive CDF and the empirical CDF of the parameter. The average predictive CDF is defined as

$$\hat{F}_I(\theta) \equiv \frac{1}{N} \sum_{i=1}^N F_i(\theta), \quad (\text{A2})$$

where N is the number of test samples and $F_i(\theta)$ is the CDF of the posterior distribution of the parameter given the data of the i th sample. The empirical CDF is defined as

$$\tilde{G}_I(\theta) \equiv \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\tilde{\theta}_i \leq \theta\}, \quad (\text{A3})$$

where the indicator function $\mathbb{1}\{A\}$ returns a value of unity if the condition A is true, and zero otherwise.

If the posteriors are marginally calibrated, $\hat{F}_I(\theta)$ and $\tilde{G}_I(\theta)$ should agree with each other for any parameter θ , so their difference should be small.

A.1.3. Hypothesis Tests

To provide a quantitative measure of the similarity between the distribution of PIT and a uniform distribution, we adopt two metrics: the Kolmogorov–Smirnov (KS; Kolmogorov 1992) test and Cramér–von Mises (CvM; Anderson 1962) test. The classical KS test is based on the distance measure D defined by

$$D \equiv \max_x |F_N(x) - F(x)|, \quad (\text{A4})$$

where $F_N(x)$ is the empirical CDF of data set x_1, x_2, \dots, x_N , and x_i is the PIT value at $\tilde{\theta}_i$ of the i th sample. $F(x)$ is the CDF of the theoretical uniform distribution. The null hypothesis of the KS test is that the variable x_i observes a uniform distribution. The p -value of the KS test is the probability to obtain a sample data distribution with D larger than the measured D_{obs} . We implement the KS test with the SciPy package under the “exact” mode (Simard & L’Ecuyer 2011). It is difficult to write down the exact formula of the p -value in the KS test. To give some intuition, it can be approximated by Ivezić et al. (2014)

$$\Pr(D > D_{\text{obs}}) = \mathcal{Q}_{\text{KS}}([\sqrt{N} + 0.12 + 0.11/\sqrt{N}] D_{\text{obs}}). \quad (\text{A5})$$

Here, the survival function \mathcal{Q}_{KS} is defined as

$$\mathcal{Q}_{\text{KS}}(u) = 2 \sum_{m=1}^{\infty} (-1)^{m-1} \exp(-2m^2 u^2). \quad (\text{A6})$$

The CvM test is based on the distance measure D^* , defined as

$$D^* \equiv N \int_{-\infty}^{\infty} (F_N(x) - F(x))^2 dF(x). \quad (\text{A7})$$

The p -value of the CvM test is the probability to obtain a sample data distribution with D^* larger than the measured D_{obs}^* , which is given (Csörgo & Faraway 1996) in the SciPy package.

The KS test and the CvM test focus on different aspects of a distribution: while the KS test is sensitive to the median, the CvM test can capture the tails of a distribution. If the p -value is greater than some criterion, typically 0.01 or 0.05 (Ivezić et al. 2014), then the null hypothesis that the distribution is a uniform distribution can be accepted.

A.2. Joint Posteriors Validation

Since parameters are degenerate, validation of the marginal posterior can be biased. The tools in this subsection focus on the validation of posteriors in the joint parameter space.

A.2.1. Probabilistic Copula Calibration

The probabilistic copula calibration (Ziegel & Gneiting 2014; Mucesh et al. 2021) is an extension of the probabilistic calibration. Consider a joint parameter distribution $f(\theta)$, where θ is a vector of parameters. For example, $\theta = \{\log_{10}(T_{\text{vir}}), \log_{10} \zeta\}$ in this paper, and $f(\theta)$ is the multidimensional probability distribution function that is generated in the MCMC chain given the observed data.

We define the Kendall distribution function

$$\mathcal{K}_H(w) \equiv \Pr\{H(\theta) \leq w\} \quad \text{for } w \in [0, 1], \quad (\text{A8})$$

where $H(\theta)$ is the CDF of the distribution $f(\theta)$, and “Pr” means the probability. Thus, the Kendall distribution function can be interpreted as the CDF of the CDF.

We define the copula probability integral transformation (copPIT) as the Kendall distribution function evaluated at the CDF of the true value $H(\tilde{\theta})$, i.e.,

$$\text{copPIT} \equiv \mathcal{K}_H(H(\tilde{\theta})) = \Pr\{H(\theta) \leq H(\tilde{\theta})\}. \quad (\text{A9})$$

The joint posteriors are probabilistically copula calibrated if the copPIT is uniformly distributed. In the QQ plot, wherein the quantile of the copPIT distribution from the data is compared with that from a uniform distribution of copPIT, this QQ plot falls on the diagonal line if the copPIT distribution is exactly uniform.

A.2.2. Kendall Calibration

As an extension of the marginal calibration, the Kendall calibration (Ziegel & Gneiting 2014; Mucesh et al. 2021) compares the average Kendall distribution function with the empirical CDF. The average Kendall distribution function is defined as

$$\hat{\mathcal{K}}_{H_i}(w) \equiv \frac{1}{N} \sum_{i=1}^N \mathcal{K}_{H_i}(w), \quad (\text{A10})$$

where $\mathcal{K}_{H_i}(w)$ is the Kendall distribution function of the CDF $H_i(\theta)$ for the i th sample. The empirical CDF of the joint CDFs evaluated at the true values $\tilde{\theta}$ is defined as

$$\tilde{J}_i(w) \equiv \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{H_i(\tilde{\theta}_i) \leq w\}. \quad (\text{A11})$$

Kendall calibration tests whether $\hat{\mathcal{K}}_{H_i}(w) - \tilde{J}_i(w)$ is small for any value $0 \leq w \leq 1$. In other words, Kendall calibration probes how well the inferred distribution agrees with the real distribution on average.

A.2.3. Probabilistic HPD Calibration

The highest probability density (HPD) is defined as (Harrison et al. 2015)

$$\text{HPD}(\tilde{\theta}) \equiv \int_{f(\theta) \geq f(\tilde{\theta})} f(\theta) d^n \theta, \quad (\text{A12})$$

where $d^n \theta$ is a volume element in the multidimensional parameter space.

The HPD describes the plausibility of $\tilde{\theta}$ under the distribution $f(\theta)$. A small value indicates high plausibility. The joint posteriors are probabilistically HPD calibrated, if the HPD is uniformly distributed.

Note that, for the PIT, copPIT, and HPD, the uniformity of their distributions is a necessary, but not sufficient, condition for accurate posteriors (Gneiting et al. 2007; Ziegel & Gneiting 2014; Harrison et al. 2015; Zhao et al. 2021).

A.2.4. Hypothesis Tests

The data x , or the PIT value, in Equations (A4) and (A7) can be replaced by the copPIT value and the HPD value, in the KS test and CvM test, respectively.

Appendix B Effect of the Sample Size

In mock observations with thermal noise and foreground cut, 9000 samples were employed for training in this paper. However, the size of this database might be larger than




Figure 8. The effect of the training sample size on the posterior inference. We show the p -values of the KS test (solid line) and CvM test (dashed line) for the null hypothesis that a statistic is of a uniform distribution, as a function of the training sample size. The tests are performed over 300 testing samples for which the mock observations are applied with thermal noise of HERA (left) and SKA (right), respectively. The statistical metrics include the PIT for $\log_{10}(T_{\text{vir}})$ (red dots) and $\log_{10}\zeta$ (red squares), the copPIT (cyan dots), and HPD (blue dots). The threshold of the p -value above which the null hypothesis is accepted is marked (black solid line).

necessary. In this section, we investigate the effect of the training sample size on the posterior inference.

The p -value of the null hypothesis that a statistic is of a uniform distribution is a good indicator for validating the inferred posteriors. We adopt the criterion that the significance level is above 0.01 for accepting the null hypothesis. The p -value is typically affected by the sample size, as well as the complexity of the individual NDE and the ensembles of multiple NDEs. With the chosen NDE architecture and an ensemble of five such NDEs, we use the bisection method and decrease the sample size for training the NDEs to find out the minimum size of training samples that meets the criterion of accepting the null hypothesis.

In Figure 8, we plot the p -values for the null hypotheses of the statistics as a function of the training sample size. For the mock HERA (SKA) observations, the p -values for all statistics are above 0.01 when the sample size is larger than about 1500 (1000). This test indicates that the minimum sample size for training can be about 1500, in general. However, since the hypothesis tests are only necessary, not sufficient, conditions for accurate posteriors, a training data set of only 1500 samples may not be optimal. Also, our estimate of the minimum sample size is only based on a two-parameter space of reionization model. The scaling law in a higher-dimensional parameter space is beyond the scope of this paper.

ORCID iDs

Xiaosheng Zhao <https://orcid.org/0000-0002-8328-1447>

Yi Mao <https://orcid.org/0000-0002-1301-3893>

Benjamin D. Wandelt <https://orcid.org/0000-0002-5854-8269>

References

Abadi, M., Barham, P., Chen, J., et al. 2016, arXiv:1605.08695
 Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *ApJ*, **925**, 221
 Alsing, J., Charnock, T., Feeney, S., & Wandelt, B. 2019, *MNRAS*, **488**, 4440
 Alsing, J., Wandelt, B., & Feeney, S. 2018, *MNRAS*, **477**, 2874
 Anderson, T. W. 1962, *Ann. Math. Statist.*, **33**, 1148
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Beardsley, A. P., Morales, M. F., Lidz, A., Malloy, M., & Sutter, P. M. 2015, *ApJ*, **800**, 128
 Bishop, C. M. 1994, Mixture Density Networks, Technical Report NCRG/94/004, Birmingham Aston Univ. <https://publications.aston.ac.uk/id/eprint/373/>
 Cameron, E., & Pettitt, A. N. 2012, *MNRAS*, **425**, 44
 Choudhury, M., Datta, A., & Majumdar, S. 2022, *MNRAS*, **512**, 5010
 Cranmer, K., Brehmer, J., & Louppe, G. 2020, *PNAS*, **117**, 30055
 Csőrgő, S., & Faraway, S. 1996, *J. R. Statist. Soc. B*, **58**, 221
 Dayal, P., & Ferrara, A. 2018, *PhR*, **780**, 1
 DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
 Doussot, A., Eames, E., & Semelin, B. 2019, *MNRAS*, **490**, 371
 Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *PhR*, **433**, 181
 Gillet, N., Mesinger, A., Greig, B., Liu, A., & Ucci, G. 2019, *MNRAS*, **484**, 282
 Gneiting, T., Balabdaoui, F., & Raftery, A. E. 2007, *J. R. Statist. Soc. B*, **69**, 243
 Greig, B., & Mesinger, A. 2015, *MNRAS*, **449**, 4246
 Greig, B., & Mesinger, A. 2017, *MNRAS*, **472**, 2651
 Greig, B., & Mesinger, A. 2018, *MNRAS*, **477**, 3217
 Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017, *MNRAS*, **466**, 4239
 Hahn, C., Vakili, M., Walsh, K., et al. 2017, *MNRAS*, **469**, 2791
 Hamill, T. M. 2001, *MWRv*, **129**, 550
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Harrison, D., Sutton, D., Carvalho, P., & Hobson, M. 2015, *MNRAS*, **451**, 2610
 Hassan, S., Andrianomena, S., & Doughty, C. 2020, *MNRAS*, **494**, 5761
 Hermans, J., Delaunoy, A., Rozet, F., Wehenkel, A., & Louppe, G. 2021, arXiv:2110.06581
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, **598**, A78
 Ivezić, Z., Connolly, A. J., VanderPlas, J. T., & Gray, A. 2014, Statistics, Data Mining, and Machine Learning in Astronomy: A Practical Python Guide for the Analysis of Survey Data (Princeton, NJ: Princeton Univ. Press)
 Jennings, W. D., Watkinson, C. A., Abdalla, F. B., & McEwen, J. D. 2019, *MNRAS*, **483**, 2907
 Kannan, R., Garaldi, E., Smith, A., et al. 2022, *MNRAS*, **511**, 4005
 Kern, N. S., Liu, A., Parsons, A. R., Mesinger, A., & Greig, B. 2017, *ApJ*, **848**, 23
 Kolmogorov, A. 1992, in On the Empirical Determination of a Distribution Function, ed. S. Kotz & N. L. Johnson (New York: Springer), 106
 Lewis, A. 2019, arXiv:1910.13970
 Mao, Y., Shapiro, P. R., Mellema, G., et al. 2012, *MNRAS*, **422**, 926
 Mellema, G., Koopmans, L. V. E., Abdalla, F. A., et al. 2013, *ExA*, **36**, 235

- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2020, *MNRAS*, **493**, 1662
- Mesinger, A., & Furlanetto, S. 2007, *ApJ*, **669**, 663
- Mesinger, A., Furlanetto, S., & Cen, R. 2011, *MNRAS*, **411**, 955
- Mondal, R., Bharadwaj, S., & Majumdar, S. 2016, *MNRAS*, **464**, 2992
- Mucesh, S., Hartley, W. G., Palmese, A., et al. 2021, *MNRAS*, **502**, 2770
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, **433**, 639
- Papamakarios, G. 2019, arXiv:1910.13233
- Papamakarios, G., Pavlakou, T., & Murray, I. 2017, in Proc. 31st Int. Conf. on Neural Information Processing Systems, NIPS'17, ed. I. Guyon et al. (Red Hook, NY: Curran Associates Inc.), 2335
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *ApJ*, **139**, 1468
- Parsons, A. R., Liu, A., Aguirre, J. E., et al. 2014, *ApJ*, **788**, 106
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825, <http://jmlr.org/papers/v12/pedregosa11a.html>
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13
- Pober, J. C., Parsons, A. R., DeBoer, D. R., et al. 2013, *AJ*, **145**, 65
- Pober, J. C., Liu, A., Dillon, J. S., et al. 2014, *ApJ*, **782**, 66
- Pober, J. C., Ali, Z. S., Parsons, A. R., et al. 2015, *ApJ*, **809**, 62
- Prelogović, D., Mesinger, A., Murray, S., Fiameni, G., & Gillet, N. 2022, *MNRAS*, **509**, 3852
- Schafer, C. M., & Freeman, P. E. 2012, in *Statistical Challenges in Modern Astronomy V*, ed. E. D. Feigelson & G. J. Babu (New York: Springer), 3
- Schmit, C. J., & Pritchard, J. R. 2018, *MNRAS*, **475**, 1213
- Shaw, A. K., Bharadwaj, S., & Mondal, R. 2019, *MNRAS*, **487**, 4951
- Shaw, A. K., Bharadwaj, S., & Mondal, R. 2020, *MNRAS*, **498**, 1480
- Shimabukuro, H., & Semelin, B. 2017, *MNRAS*, **468**, 3869
- Sikder, S., Barkana, R., Reis, I., & Fialkov, A. 2022, arXiv:2201.08205
- Simard, R., & L'Ecuyer, P. 2011, *J. Stat. Softw.*, **39**, 1
- Thompson, A. R., Moran, J. M., & Swenson, G. W. J. 2001, *Interferometry and Synthesis in Radio Astronomy* (2nd ed.; New York, NY: Wiley)
- Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, *PASA*, **30**, e007
- Trott, C. M., Jordan, C. H., Midgley, S., et al. 2020, *MNRAS*, **493**, 4711
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, **556**, A2
- Van Rossum, G., & Drake, F. L. 2009, *Python 3 Reference Manual* (Scotts Valley, CA: CreateSpace)
- Van Rossum, G., & Drake, F. L., Jr 1995, *Python Reference Manual* (Amsterdam: Centrum voor Wiskunde en Informatica Amsterdam)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Waskom, M. L. 2021, *JOSS*, **6**, 3021
- Watkinson, C. A., Greig, B., & Mesinger, A. 2022, *MNRAS*, **510**, 3838
- Zhao, D., Dalmasso, N., Izbicki, R., & Lee, A. B. 2021, in Proc. 37th Conf. on Uncertainty in Artificial Intelligence, 161, ed. C. de Campos & M. H. Maathuis (Online: PMLR), 1830, <https://proceedings.mlr.press/v161/zhao21b.html>
- Zhao, X., Mao, Y., Cheng, C., & Wandelt, B. D. 2022, *ApJ*, **926**, 151
- Zhou, Y., & La Plante, P. 2022, *PASP*, **134**, 044001
- Ziegel, J. F., & Gneiting, T. 2014, *Electron. J. Statist.*, **8**, 2619