



Exploring the likelihood of the 21-cm power spectrum with simulation-based inference

David Prelogović * and Andrei Mesinger 

Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56125 Pisa, Italy

Accepted 2023 July 3. Received 2023 June 25; in original form 2023 May 4

ABSTRACT

Observations of the cosmic 21-cm power spectrum (PS) are starting to enable precision Bayesian inference of galaxy properties and physical cosmology, during the first billion years of our Universe. Here we investigate the impact of common approximations about the likelihood used in such inferences, including: (i) assuming a Gaussian functional form; (ii) estimating the mean from a single realization; and (iii) estimating the (co)variance at a single point in parameter space. We compare ‘classical’ inference that uses an explicit likelihood with simulation-based inference (SBI) that estimates the likelihood from a training set. Our forward models include: (i) realizations of the cosmic 21-cm signal computed with 21CMFAST by varying ultraviolet (UV) and X-ray galaxy parameters together with the initial conditions; (ii) realizations of the telescope noise corresponding to a 1000 h integration with the low-frequency component of the Square Kilometre Array (SKA1-Low); and (iii) the excision of Fourier modes corresponding to a foreground-dominated horizon ‘wedge’. We find that the 1D PS likelihood is well described by a Gaussian accounting for covariances between wave modes and redshift bins (higher order correlations are small). However, common approaches of estimating the forward-modelled mean and (co)variance from a random realization or at a single point in parameter space result in biased and overconstrained posteriors. Our best results come from using SBI to fit a non-Gaussian likelihood with a Gaussian mixture neural density estimator. Such SBI can be performed with up to an order of magnitude fewer simulations than classical, explicit likelihood inference. Thus SBI provides accurate posteriors at a comparably low computational cost.

Key words: methods: data analysis – methods: statistical – dark ages, reionization, first stars – cosmology: theory.

1 INTRODUCTION

The redshifted 21-cm signal, corresponding to the hyperfine splitting of the 1S ground state of neutral hydrogen, is becoming a powerful probe of the first billion years of our Universe. This evolutionary milestone includes the Cosmic Dawn (CD) of the first galaxies and the eventual reionization of the intergalactic medium (IGM). Containing orders of magnitude more independent modes than the cosmic microwave background (CMB; e.g. Loeb & Zaldarriaga 2004), this cosmic 21-cm signal will allow us to infer the properties of the unseen first galaxies (e.g. Park et al. 2019), together with physical cosmology (e.g. Kern et al. 2017). Current radio telescopes are already putting compelling upper limits on its spherically averaged power spectrum (PS; e.g. Mertens et al. 2020; Trott et al. 2020; Abdurashidova et al. 2022a,b). The upcoming Square Kilometre Array (SKA)¹ should not only *detect* the PS, but also enable image-space mapping of the signal (Koopmans et al. 2015; Mesinger 2020).

There are several important reasons why these experiments primarily aim to detect the PS (as compared with other summary statistics). First, it is the most physically motivated summary: the theoretical PS can be analytically computed for simple functions of the matter field that is Gaussian on large scales (e.g. Barkana & Loeb

2005; Pritchard & Furlanetto 2007; McQuinn & D’Aloisio 2018; Schneider, Giri & Mirocha 2021; Muñoz 2023). Secondly, due to its high compression (many Fourier modes are sampled to compute the PS in a given bin), it provides a good signal-to-noise ratio (S/N) facilitating a detection with limited, preliminary data. Finally, the visibility space of interferometric measurements is closely related to Fourier space. This means that the telescope noise, systematics, and foreground contamination can be more naturally characterized in PS space, compared with e.g. image-based summaries. As a result, inference from the 21-cm signal generally uses the PS when computing a likelihood (though see e.g. Shimabukuro et al. 2017; Gazagnes, Koopmans & Wilkinson 2021; Watkinson, Greig & Mesinger 2022; Greig, Ting & Kaurov 2023).

What is the form of the likelihood of the 21-cm PS? Virtually all works (e.g. Greig & Mesinger 2015; Ghara et al. 2020; Greig et al. 2021; HERA Collaboration et al. 2023; Maity & Choudhury 2023) assume a Gaussian form for the likelihood:

$$\begin{aligned} \ln \mathcal{L}(\Delta_{21\text{obs}}^2 | \boldsymbol{\theta}) &\propto -\frac{1}{2} [\Delta_{21\text{obs}}^2 - \boldsymbol{\mu}(\boldsymbol{\theta})]^T \boldsymbol{\Sigma}^{-1} [\Delta_{21\text{obs}}^2 - \boldsymbol{\mu}(\boldsymbol{\theta})] \\ &\approx -\frac{1}{2} \sum_{k,z} \frac{[\Delta_{21\text{obs}}^2(k, z) - \boldsymbol{\mu}(\boldsymbol{\theta}, k, z)]^2}{\sigma^2(\boldsymbol{\theta}, k, z)}, \end{aligned} \quad (1)$$

where $\Delta_{21\text{obs}}^2(k, z)$ is the ‘observed’ PS (cf. equation 4). For a given astrophysical/cosmological parameter vector $\boldsymbol{\theta}$, the forward-modelled PS is *stochastic*, primarily due to cosmic variance and

* E-mail: david.prelogovic@sns.it

¹<https://www.skatelescope.org>

thermal noise: $\mu(\boldsymbol{\theta}, k, z)$ is the mean of this distribution and $\Sigma(\boldsymbol{\theta}, k, z, k', z')$ its covariance matrix capturing two-point correlations between wave modes and redshifts (note we omit the explicit summation over bins, writing it as a vector–matrix product for simplicity). Most often, the full covariance matrix is approximated to be diagonal with variance $\sigma^2(\boldsymbol{\theta}, k, z)$ (though see Nasirudin et al. 2020), leading to the simple summation over (k, z) bins. Furthermore, due to computational requirements, the forward-modelled mean and (co)variance are not computed on-the-fly at each $\boldsymbol{\theta}$ by varying the cosmic seed and other sources of stochasticity. Instead, a single realization is used in place of the mean, and the (co)variance is generally only computed at a single, fiducial parameter, $\boldsymbol{\theta}_{\text{fid}}$, and then kept fixed during inference (though see Watkinson et al. 2022). One then infers the posterior probability distribution $P(\boldsymbol{\theta}|\Delta_{21\text{obs}}^2)$, using the Bayes’ equation:

$$P(\boldsymbol{\theta}|\Delta_{21\text{obs}}^2) = \frac{\mathcal{L}(\Delta_{21\text{obs}}^2|\boldsymbol{\theta}) \cdot \pi(\boldsymbol{\theta})}{\mathcal{Z}(\Delta_{21\text{obs}}^2)}, \quad (2)$$

where $\pi(\boldsymbol{\theta})$ the prior probability of $\boldsymbol{\theta}$, and the evidence $\mathcal{Z}(\Delta_{21\text{obs}}^2)$ is here treated as a normalization constant as we are focusing only on a single model.

The form of the likelihood in equation (1) is only guaranteed to be true for a Gaussian random field. Although the primordial density fluctuations are indeed Gaussian, subsequent non-linear evolution through gravitational collapse and radiative cooling introduces non-Gaussianities in the hydrogen density field on moderate to small scales. More importantly, the cosmic 21-cm signal is also sensitive to the temperature and ionization state of the IGM, which is in turn driven by the multiwavelength radiation fields of the first galaxies. Indeed, this dependence is precisely what allows us to infer the properties of these galaxies; however, it introduces highly non-trivial mode coupling across a wide range of scales. As a result, the redshifted 21-cm signal from the CD is *not* Gaussian (e.g. Mellema et al. 2013; Greig, Ting & Kaurov 2022; Watkinson et al. 2022). This has two important implications. First, the PS does not provide optimal compression; other summary statistics could in principle better constrain astrophysical and cosmological parameters (e.g. Gazagnes et al. 2021; Zhao et al. 2022a; Greig et al. 2023; Prelogović & Mesinger, in preparation). Secondly, the likelihood of the PS is *not* Gaussian distributed according to equation (1). Quantifying this later point is the subject of this work.

How can one determine the form of the 21-cm PS likelihood? Ideally, one would start with the Gaussian initial conditions and trace the likelihood throughout the evolution of the Universe, either analytically or numerically. Indeed, the field of large-scale structure (LSS) as probed by galaxy surveys has made impressive progress along those lines (e.g. Kitaura & Enßlin 2008; Jasche & Kitaura 2010; Jasche & Wandelt 2012, 2013; Leclercq et al. 2017; Dai & Seljak 2022; McAlpine et al. 2022; Bayer, Modi & Ferraro 2023). However, the dependence on unknown, multiscale, multiwavelength radiation fields makes it more difficult to obtain a tractable likelihood for the 21-cm PS.

Another option is to use simulation-based inference (SBI; for a recent review see Cranmer, Brehmer & Loupé 2020). SBI relies on a training set of simulated data, which is used to train neural density estimators (NDEs) to fit the likelihood without having to specify its functional form (indeed SBI is commonly also referred to as ‘likelihood-free’ or ‘implicit likelihood’ inference). Specifically, SBI consists of the following general steps.

- (i) Sample a parameter vector from the prior $\tilde{\boldsymbol{\theta}} \sim \pi(\boldsymbol{\theta})$.

- (ii) Simulate a data vector corresponding to the parameter sample, $\tilde{\boldsymbol{d}} = \text{simulator}(\tilde{\boldsymbol{\theta}})$. This is equivalent to drawing from the likelihood function $\tilde{\boldsymbol{d}} \sim \mathcal{L}(\boldsymbol{d}|\boldsymbol{\theta})$.

- (iii) Repeat many times.

Here the data space \boldsymbol{d} corresponds to the 21-cm PS $\Delta_{21}^2(k, z)$. For the simulator of the 21-cm signal, here we use the public 21CMFAST code, described in more detail in Section 2. The resulting set $\{(\tilde{\boldsymbol{d}}_1, \tilde{\boldsymbol{\theta}}_1), \dots, (\tilde{\boldsymbol{d}}_N, \tilde{\boldsymbol{\theta}}_N)\}$ corresponds to samples from the joint distribution $P(\boldsymbol{d}, \boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{d}|\boldsymbol{\theta}) \cdot \pi(\boldsymbol{\theta})$. One can then use a density estimator to fit the likelihood $\mathcal{L}(\boldsymbol{d}|\boldsymbol{\theta})$. Note that there are many variants to such a procedure (e.g. Cranmer et al. 2020), including fitting the posterior directly or fitting likelihood ratios (e.g. Cole et al. 2022). SBI is rapidly becoming popular in the field of cosmic 21-cm due to its ability to deliver efficient, unbiased inference (Zhao et al. 2022a; Zhao, Mao & Wandelt 2022b; Saxena et al. 2023).

In this work, we use SBI to test the validity of a wide range of approximations for the likelihood of the spherically averaged 21-cm PS. These include: (i) a 1D variance evaluated at a single parameter choice; (ii) a 2D covariance evaluated at a single parameter choice; (iii) a 1D variance evaluated on-the-fly; (iv) a 2D covariance evaluated on-the-fly; and (v) a fully non-Gaussian likelihood (allowing us to capture all higher order moments between k and z bins). We compare approaches on the basis of both accuracy (verified in some cases using simulation-based calibration – SBC) and computational efficiency. As this work was nearing completion, a similar study was published by Zhao et al. (2022b), comparing SBI to explicit likelihood inference in the recovery of two astrophysical parameters from the 21-cm PS. Our work is an improvement because: (i) we use a larger astrophysical parameter space; (ii) we compute all summaries on the light-cone (instead of computing some on light-cones and others on coeval cubes), thus allowing for a more precise comparison between different likelihood estimators; (iii) we compute the noise more accurately based on baseline sampling directly in uv space; and (iv) we explore a broader set of functional forms for the likelihood.

This paper is organized as follows. In Section 2, we discuss the main ingredients of the inference, including our forward simulator, 21CMFAST, astrophysical parameters, and priors. In Section 3, we discuss the various functional forms of the PS likelihood we explore and the NDEs used to fit them. In Section 4, we show results for a fiducial mock observation, and test the inference across our parameter space with SBC. Finally, we conclude in Section 5. All quantities are quoted in comoving units, and we assume a standard Λ cold dark matter (Λ CDM) cosmology: $(\Omega_\Lambda, \Omega_M, \Omega_b, n, \sigma_8, H_0) = (0.69, 0.31, 0.048, 0.97, 0.81, 68 \text{ km s}^{-1} \text{ Mpc}^{-1})$, consistent with the results from Planck Collaboration VI (2020).

2 SIMULATING 21-CM OBSERVATIONS

Inference, whether using SBI or with an explicit likelihood, requires an accurate simulator to generate mock observables from samples of astrophysical/cosmological parameters. Our simulation pipeline is illustrated in Fig. 1. It consists of the following steps.

- (i) *Cosmological signal.* We simulate a realization of the 21-cm light-cone (more precisely a light ‘cuboid’), corresponding to a sampled parameter vector, $\tilde{\boldsymbol{\theta}}$, and a sampled random seed for generating the initial conditions.

- (ii) *Mean removal.* We remove the mean from each redshift/frequency slice, to account for the inability of interferometers to measure the $k_\perp = 0$ mode.

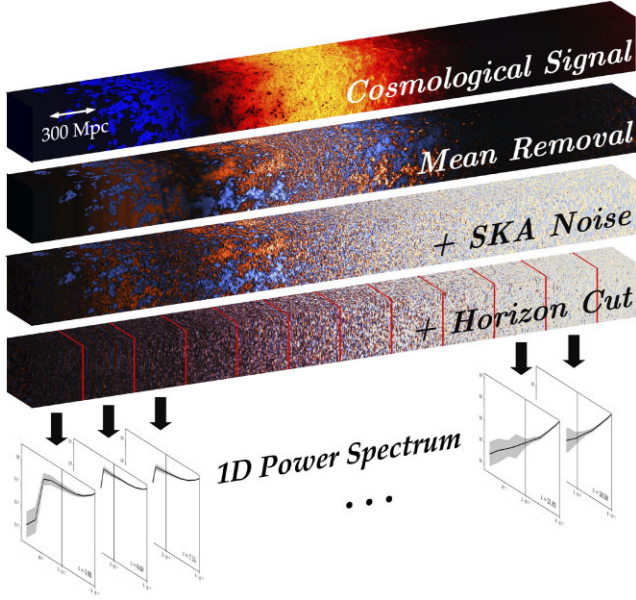


Figure 1. Schematic of our simulation pipeline. Starting from the cosmological signal computed with 21CMFAST, we remove the mean of the signal, add noise corresponding to a 1000 h SKA1-Low observation, and perform a foreground cut below the horizon limit (see text for details). Finally, we bin the light-cone and compute the 1D power spectrum (PS) in each bin. For the reference, in the upper left-hand corner of the first light-cone, we show the scale of the simulation box size.

(iii) *+SKA noise.* We add a realization of noise corresponding to a 1000 h integration with the low-frequency component of the Square Kilometre Array (SKA1-Low).

(iv) *+Horizon cut.* We remove a foreground-dominated ‘wedge’ region by zeroing the corresponding Fourier modes.

(v) *1D PS.* We cut the light-cone into blocks of equal conformal length along the redshift axis, computing the 1D PS for each block. This results in $\Delta_{21}^2(k, z)$ that we use as our summary statistics throughout this paper.

Below we briefly describe these steps in more detail.

To compute the cosmological signal from the first step, we use the public, seminumerical code 21CMFASTv3² (Mesinger, Furlanetto & Cen 2011; Murray et al. 2020), with the galaxy parametrization from Park et al. (2019). The code generates a realization of the initial density and velocity fields, and evolves them with second-order Lagrangian perturbation theory (2LPT; Scoccimarro 1998). From the evolved density fields, the conditional halo mass function is used to compute the spatial fluctuations in the galaxy field (e.g. Barkana & Loeb 2004). The inhomogeneous reionization field is obtained by comparing the number of ionizing photons to the number of recombinations, in regions of decreasing radii (Furlanetto, Zaldarriaga & Hernquist 2004; Sobacchi & Mesinger 2014). Soft ultraviolet (UV) and X-ray photons that have much longer mean free paths are instead tracked by integrating the local emissivity back along the light-cone, for each simulation cell. These radiation fields then impact the temperature and ionization state of each IGM cell. For more details, interested readers are encouraged to see Mesinger & Furlanetto (2007) and Mesinger et al. (2011).

Our simulations correspond to 300 Mpc boxes, with a cell size of 1.5 Mpc. The choices of astrophysical galaxy parameters are

discussed in the following section. We interpolate between adjacent comoving snapshots, also accounting for subgrid redshift-space distortions (e.g. Mao et al. 2012; Jensen et al. 2013; Greig & Mesinger 2018), creating a light-cone of the cosmic signal extending from redshift 30 to 5 (see the top panel of Fig. 1).

SKA1-Low uv coverage and thermal noise are calculated using TOOLS21CM³ (Giri, Mellema & Jensen 2020). We assume a tracked scan of 6 h d⁻¹, 10 s integration time, for a total of 1000 h, using only the core stations (baseline ≤ 2 km). After subtracting the mean signal from each slice in the light-cone and adding the thermal noise corresponding to this uv coverage, we also remove a foreground-contaminated ‘wedge’ (Morales et al. 2012; Trott, Wayth & Tingay 2012; Vedantham, Udaya Shankar & Subrahmanyan 2012; Liu, Parsons & Trott 2014a,b; Parsons et al. 2014; Pober et al. 2014; Murray & Trott 2018; Liu & Shaw 2020). Conservatively, we remove (zero) all modes below the horizon limit, which can be expressed as a slope in the line-of-sight (k_{\parallel}) versus sky plane (k_{\perp}) Fourier modes:

$$k_{\parallel} \leq k_{\perp} \frac{E(z)}{1+z} \int_0^z \frac{dz'}{E(z')}, \quad (3)$$

where $E(z) = \sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$. For more details on the telescope noise and foreground avoidance implemented in this work, we refer reader to Prelogović et al. (2022). The only difference with respect to the original method is that we do not apply a ‘rolling’ of the wedge filter, more relevant for 21-cm images. As the 1D PS used here is computed in (binned) Fourier space, it is sufficient to apply the wedge filter once per the light-cone chunk on which the PS is calculated.

After the telescope effects are included, we cut the resulting 21-cm brightness temperature light-cone, $\delta T_b(\mathbf{x}, z)$, along redshift axis into chunks of 300 Mpc and compute the 1D PS on each section as

$$\delta \bar{T}_b^2 \Delta_{21}^2(k, z) \equiv \frac{k^3}{2\pi^2 V} \left\langle |\delta T_b(\mathbf{k}, z) - \delta \bar{T}_b(z)|^2 \right\rangle_k, \quad (4)$$

where z is the central redshift of each chunk. This 1D PS serves as our summary statistic throughout this work.

2.1 Model parameters

We characterize the unknown UV and X-ray properties of high- z galaxies using the model from Park et al. (2019). The simple scaling relations from that work are both easy to interpret (thus allowing us to set well-motivated priors) and sufficient to recover currently available galaxy and Epoch of Reionization (EoR) observations. Specifically, we use a subset of five parameters from Park et al. (2019) that drives most of the variation in the signal.

(i) $f_{*,10}$ – fraction of the galactic gas in stars, normalized at a halo mass of $10^{10} M_\odot$. The stellar mass of a galaxy M_* is assumed to be proportional to the halo mass M_h as

$$M_* = f_* \left(\frac{\Omega_b}{\Omega_m} \right) M_h,$$

where f_* is a power law,

$$f_*(M_h) = f_{*,10} \left(\frac{M_h}{10^{10} M_\odot} \right)^{\alpha_*}.$$

²<https://github.com/21cmfast/21cmFAST>

³<https://github.com/sambit-giri/tools21cm>

Here, we fix the power index to $\alpha_* = 0.5$, which fits the observed, faint-end⁴ UV luminosity functions very well (see Park et al. 2019 and references therein).

(ii) $f_{\text{esc},10}$ – the ionizing UV escape fraction, normalized at a halo mass of $10^{10} M_\odot$. The ionizing escape fraction is taken to scale with the halo mass as a power law,

$$f_{\text{esc}}(M_h) = f_{\text{esc},10} \left(\frac{M_h}{10^{10} M_\odot} \right)^{\alpha_{\text{esc}}}.$$

Here again we fix the power-law index to $\alpha_{\text{esc}} = -0.5$ (cf. Park et al. 2019).

(iii) M_{turn} – a characteristic halo mass below which star formation becomes inefficient due to slow gas accretion, photoheating, and/or supernovae feedback. We assume only a fraction of haloes given by

$$f_{\text{duty}}(M_h) = \exp\left(-\frac{M_{\text{turn}}}{M_h}\right)$$

can host star-forming galaxies (cf. Sobacchi & Mesinger 2013a,b; O’Shea et al. 2015; Xu et al. 2016).

(iv) E_0 – the energy threshold below which X-ray photons are absorbed by the interstellar medium (ISM) of the host galaxy. This will depend on the ISM density, metallicity, and the local environment of X-ray sources (like high-mass X-ray binary stars; e.g. Das et al. 2017).

(v) $L_{X<2\text{keV}}/\text{SFR}$ – the soft-band (with energies between E_0 and 2 keV) X-ray luminosity per unit star formation escaping the galaxy, in units of $\text{erg s}^{-1} \text{keV}^{-1} M_\odot^{-1} \text{yr}$. Photons with energies greater than 2 keV have mean free paths larger than the Hubble length during the CD (e.g. McQuinn & O’Leary 2012), while we assume photons with energies below E_0 are absorbed in the host galaxy. We assume a star formation rate (SFR) of $\dot{M}_* = M_*/(t_* H^{-1})$, where t_* controls the characteristic time-scale as a fraction of Hubble time, $H(z)$. Here, we fix it to $t_* = 0.5$ (cf. Park et al. 2019).

Our data bases are created by sampling physically motivated priors for the above five galaxy parameters. We take flat priors for X-ray parameters $E_0/[\text{keV}] \in [0.2, 1.5]$ and $\log_{10} \frac{L_{X<2\text{keV}}/\text{SFR}}{[\text{erg s}^{-1} \text{keV}^{-1} M_\odot^{-1} \text{yr}]} \in [38, 42]$, as they are only weakly and indirectly constrained for the high-mass X-ray binaries that likely dominate X-ray production in the first galaxies (e.g. Mineo, Gilfanov & Sunyaev 2012; Fragos et al. 2013; Lehmer et al. 2016; Das et al. 2017). Our priors for star formation and UV emission are taken from the posterior of Park et al. (2019), constructed from existing measurements of (i) UV luminosity functions (Bouwens et al. 2015a, 2017; Oesch et al. 2018), (ii) τ_e (Planck Collaboration XLVII 2016), and (iii) the dark fraction from the Lyman forest (McGreer, Mesinger & D’Odorico 2015). Specifically, we fit the 6D posterior distribution from Park et al. (2019) with a kernel density estimator, and evaluate it at the fiducial values of ($\alpha_* = 0.5$, $t_* = 0.5$, and $\alpha_{\text{esc}} = -0.5$) to obtain our prior over $f_{*,10}$, $f_{\text{esc},10}$, and M_{turn} . For this we use the CONDITIONAL_KDE⁵ code. We stress that the posteriors that we obtain from the 21-cm PS are much narrower than the priors, which means that the inference

⁴The general stellar to halo mass relation is better fit with a double power law, with a flatter index for haloes above $\sim 10^{12} M_\odot$ (e.g. Behroozi et al. 2019; Mirocha & Furlanetto 2019). However, the radiation fields that drive the 21-cm signal are dominated by the far more abundant faint galaxies (e.g. Bouwens et al. 2015b; Nikolić et al. 2023); thus we do not bother to characterize the rare massive galaxies (or active galactic nuclei) that have a negligible contribution to cosmic radiation fields.

⁵https://github.com/dprelogo/conditional_kde

is likelihood dominated and not sensitive to our specific choice of priors.

When required below, we assume a fiducial parameter set, θ_{fid} , consistent with Park et al. (2019): $\log_{10} f_{*,10} = -1.3$, $\log_{10} f_{\text{esc},10} = -1.0$, $\log_{10} M_{\text{turn}}/M_\odot = 8.7$, $E_0 = 0.5$ keV, and $\log_{10} \frac{L_{X<2\text{keV}}/\text{SFR}}{[\text{erg s}^{-1} \text{keV}^{-1} M_\odot^{-1} \text{yr}]} = 40$. For ease of notation, in the remainder of the paper we drop the \log_{10} , units and limits, denoting the parameters as $f_{*,10}$, $f_{\text{esc},10}$, E_0 , M_{turn} , and L_X/SFR .

3 CHOOSING A LIKELIHOOD FUNCTION

In this work, our summary statistic (i.e. our ‘data space’) is the 1D PS from equation (4). Here we compare different choices for the likelihood (both explicit and implicit).

3.1 Classical (explicit) Gaussian likelihoods

By far the most common choice for the 21-cm PS likelihood is a Gaussian with a fixed variance (i.e. a diagonal covariance matrix in (k, z) space) computed at some fiducial value θ_{fid} , as in equation (1). When doing a 21-cm forecast, this θ_{fid} is chosen a priori to make a mock observation. More generally, one can use a maximum likelihood estimate or a maximum a posteriori (MAP) from inference using other data sets to choose θ_{fid} .

Here we discuss all of the ingredients for computing a Gaussian likelihood. For every point in the parameter space θ , we label the Gaussian mean, variance, or full covariance as $\mu(\theta)$, $\sigma^2(\theta)$, or $\Sigma(\theta)$, respectively. In the case of the fiducial parameters θ_{fid} , we calculate them in the following way:

$$\mu_{k_i, z_j}(\theta_{\text{fid}}) = \frac{1}{N} \sum_{n=1}^N \Delta_{21}^2(\theta_{\text{fid}}, k_i, z_j)_n, \quad (5)$$

$$\begin{aligned} \Sigma_{k_i, z_j, k_l, z_m}(\theta_{\text{fid}}) &= \frac{1}{N-1} \sum_{n=1}^N (\Delta_{21}^2(\theta_{\text{fid}}, k_i, z_j)_n - \mu_{k_i, z_j}) \\ &\quad \times (\Delta_{21}^2(\theta_{\text{fid}}, k_l, z_m)_n - \mu_{k_l, z_m}), \end{aligned} \quad (6)$$

$$\sigma_{k_i, z_j}^2(\theta_{\text{fid}}) = \Sigma_{k_i, z_j, k_i, z_j}(\theta_{\text{fid}}), \quad (7)$$

where n represents a different realization of the cosmic random seed (i.e. different initial conditions) and k_i and z_j are Fourier wave mode and redshift bins, respectively. We will see later how to estimate these quantities using machine learning. Note that there are two main sources of stochasticity in our pipeline: the cosmological signal and the telescope noise. The telescope noise does not depend on θ and can be averaged over several realizations and pre-computed. Therefore, the total mean is estimated by $\mu(\theta) = \mu_{21\text{cm}}(\theta) + \mu_{\text{noise}}$. The contributions are separable even with the wedge excision.

Finally, we can write the Gaussian likelihood for an observed PS, $\Delta_{21\text{obs}}^2$, in standard notation as

$$\mathcal{L}(\Delta_{21\text{obs}}^2 | \theta) = \mathcal{N}(\Delta_{21\text{obs}}^2 | \mu(\theta), \Sigma(\theta)). \quad (8)$$

Here, $\Sigma(\theta)$ is the covariance and $\mu(\theta)$ the PS mean, estimated at each θ . The vast majority of 21-cm literature also ignores correlations between different redshifts and k -modes:

$$\mathcal{L}(\Delta_{21\text{obs}}^2 | \theta) = \mathcal{N}(\Delta_{21\text{obs}}^2 | \mu(\theta), \sigma^2(\theta)), \quad (9)$$

where the variance $\sigma^2(\theta)$ is simply the diagonal of the covariance (equation 7).

To correctly estimate the mean and (co)variance of the forward-modelled power spectra, one would need to run several simulations changing the cosmic seed (cf. equation 5) for each θ sample. However for computational reasons, the cosmic seed is usually not varied

when performing inference and thus *the mean is estimated from only one, fixed realization* (though techniques such as ‘pairing and fixing’ could reduce the computation time; e.g. Giri et al. 2023). Moreover, the cosmic (co)variance is generally only computed at the fiducial parameter values (used to make the mock observation), i.e. constants $\Sigma(\theta_{\text{fid}})$ and $\sigma^2(\theta_{\text{fid}})$ are used in place of the θ -dependent (co)variances in equations (8) and (9). Both of these assumptions can bias inference (Mondal, Bharadwaj & Majumdar 2017; Shaw, Bharadwaj & Mondal 2019; Watkinson et al. 2022; Greig et al. 2023).

Here when performing ‘classical’ (non-amortized) inference with the explicit likelihoods given by equations (8) and (9), we pick a *model seed* that gives a realization of the 21-cm PS that is *closest* (among 1000 realizations) to the mean at θ_{fid} , and then use that same cosmic seed when sampling all of the parameter space θ . Note that this is an improvement over the most common approach of *randomly* choosing the model seed and assuming it is a proxy for the mean. Furthermore for $\Delta_{21\text{ obs}}^2$, we create a mock observation at θ_{fid} by picking a ‘reasonable’ *data seed* that is within 1σ of the mean. Finally, for computational reasons, we compute the variance and covariance from 1000 realizations only at the fiducial values, θ_{fid} , keeping these uncertainties constant throughout parameter space.

The resulting mean PS, its variance, and the realizations corresponding to the model and data seeds are shown in Fig. 2, for the fiducial parameter set θ_{fid} . Because of the wedge excision, we consider only $k \geq 0.2 \text{ Mpc}^{-1}$ in our analysis, marked with vertical line. The upper limit $k \leq 0.6 \text{ Mpc}^{-1}$ is motivated by the drop in uv coverage as one approaches the maximum core baseline (see Prelogović et al. 2022 for more details). We see from the figure that our choice of model seed matches the mean almost perfectly, at least for this fiducial parameter value, θ_{fid} . This is a computationally inexpensive way of approximating the actual mean that should be computed on-the-fly when using a Gaussian likelihood for classical inference (cf. equation 1). Furthermore, we verify that our data seed used for the mock observation is ‘reasonable’, and does not correspond to a rare outlier.

In Fig. 3, we show the correlation matrix of the 21-cm PS computed at the fiducial parameters, $\Sigma_{k_i, z_j, k_l, z_m}(\theta_{\text{fid}})$ (cf. equation 6). We calculate it on 11 distinct light-cone chunks and four logarithmically spaced k -bins, in the range $0.2\text{--}0.6 \text{ Mpc}^{-1}$. To visualize such a 4D correlation matrix (k_i, z_j, k_l, z_m), one has to flatten it to 2D, which is shown in the top-left. There are 11 distinct redshift blocks, with 4×4 k -bins. The image is further split into two triangles. The upper right half shows correlations after *mean removal*, and the lower left half shows correlations after *mean removal + SKA noise* (cf. Fig. 1). The upper right-hand panel shows a zoom-in on the first three redshift chunks of the upper left-hand panel. The bottom row shows the same quantities but including also ‘+horizon cut’.

As expected, the diagonal terms (the variance) have the strongest correlation, and the telescope noise strongly suppresses off-diagonal correlations above $z \gtrsim 10$. However, the off-diagonal terms can be significant at $z \lesssim 10$, suggesting that it is important to include the full covariance matrix in the Gaussian likelihood for inference during the EoR. Similar conclusions have also been reached by Mondal et al. (2017) and Shaw et al. (2019) (see their figs 6 and 7, respectively), where the authors analysed correlations between k -bins at fixed redshifts.

In order to gain physical intuition about the origin of the strongest (anti)correlations, in Fig. 4, we show the redshift evolution of the PS amplitude at $k = 0.25 \text{ Mpc}^{-1}$ during the EoR (upper panel), together with the neutral fraction (lower panel). The blue (red) line shows the mean computed over several realizations that have high (low) power at $z = 6.7$. One can clearly see that those cosmic

seeds that result in high 21-cm power at $z = 6.7$ subsequently have low power at $z = 5.9$, consistent with the (anti)correlation redshift lengths seen in Fig. 3. In other words, the realizations in blue have a delayed, but subsequently more rapid EoR evolution, seen in the lower panel. This is something that is expected if the EoR had a larger contribution from rare, massive haloes that are on the exponential tail of the halo mass function. These massive haloes appear later than their lighter counterparts, but their fractional abundance subsequently increases more rapidly (e.g. Barkana & Loeb 2004). We confirm this explanation by looking at the probability density functions (PDFs) of the Lagrangian densities of the realizations in red versus those in blue. Indeed, the realizations shown in blue have a larger, high- σ tail in the density PDF: these rare, dense cells increase the relative abundance of the most massive haloes (flattening the halo mass function), resulting in a delayed and more rapid EoR evolution.

In what follows, we will be fitting the likelihood with different density estimators. For this, it helps to consider the likelihood as a function in the parameter space θ and a *probability density* in the PS *data space*, \mathbf{d}_{PS} . Using this notation, a Gaussian likelihood can be written as

$$\mathcal{L}(\mathbf{d}_{\text{PS}}|\theta) = \mathcal{N}(\mathbf{d}_{\text{PS}}|\boldsymbol{\mu}(\theta), \boldsymbol{\Sigma}(\theta)). \quad (10)$$

Therefore, $\mathcal{L}(\Delta_{21\text{ obs}}^2|\theta)$ from equations (8) and (9) is a function of θ for a *given* observation $\Delta_{21\text{ obs}}^2 \in \mathbf{d}_{\text{PS}}$, while $\mathcal{L}(\mathbf{d}_{\text{PS}}|\theta)$ is a probability density over all possible observations \mathbf{d}_{PS} .

3.2 Simulation-based (‘likelihood free’) inference with neural density estimators

In the previous section, we introduced two classic likelihood choices used in the vast majority of 21-cm inference. These explicit, Gaussian likelihoods are only approximations to the true likelihood for several reasons:

- (i) the mean of the forward-modelled power spectra is estimated from one simulation, instead of covarying the initial condition (IC) random seeds to compute $\boldsymbol{\mu}(\theta) = \langle \Delta_{21}^2(\theta) \rangle_{\text{ICs}}$ at each parameter sample θ (see equation 5);
- (ii) the (co)variance is generally computed only at θ_{fid} , instead of computing it at each parameter sample θ ;
- (iii) only one- and two-point correlations between wave modes and redshifts are included; higher order terms are ignored.

In the last point, we refer to higher order moments between wave modes and redshifts used in the PS likelihood, *not* the fact that the PS itself is a two-point statistic and so is a suboptimal summary statistic for non-Gaussian fields. We will quantify the later point in future work (Prelogović et al., in preparation). With the classical methods above, relaxing any of these simplifications seriously increases the computational costs, often making non-amortized inference impractically expensive.

Simulation-based inference (SBI) offers an efficient way to address all of the above limitations, by using a simulator to draw samples from the ‘real’ likelihood. As discussed in the Introduction, this involves sampling all of the relevant sources of stochasticity in the simulator (in our case the initial conditions and thermal noise), together with the astrophysical parameters, so as to create a training set of mock observations (i.e. a joint distribution of $P(\mathbf{d}_{\text{PS}}, \theta)$).

There are many different flavours of SBI, depending on the usage case. Here we use the Density Estimation Likelihood-Free Inference (DELFI; Alsing, Wandelt & Feeney 2018; Papamakarios, Sterratt & Murray 2019; Alsing et al. 2019), which uses a density

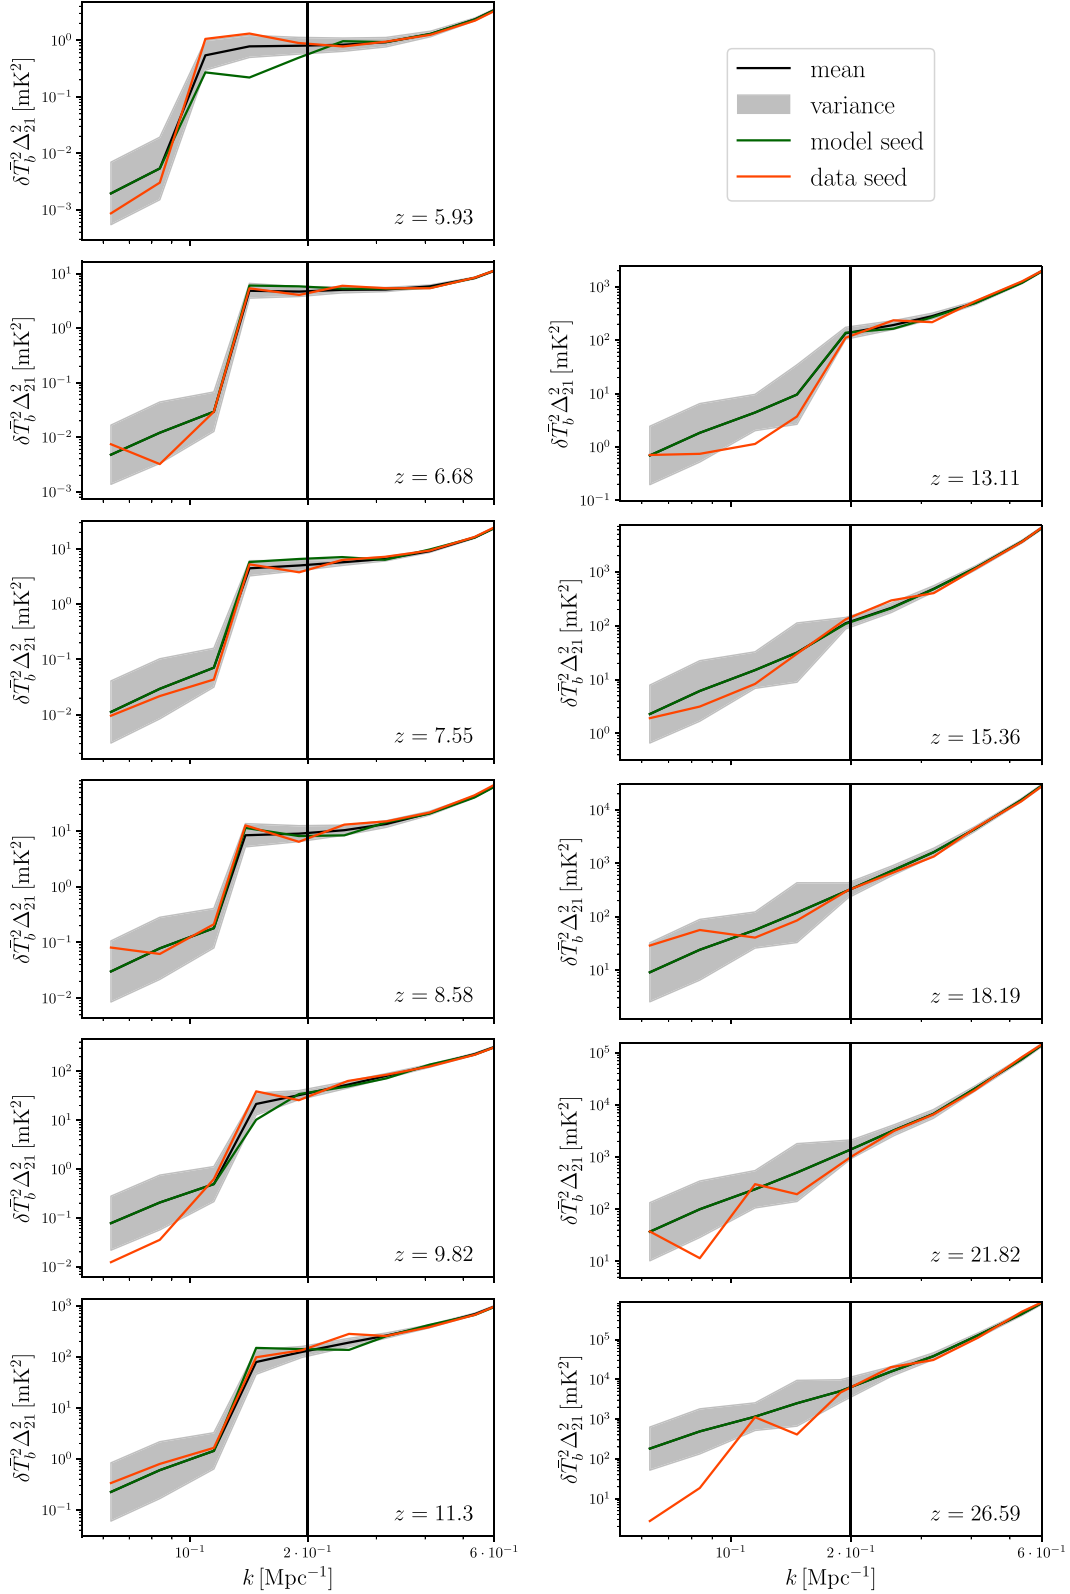


Figure 2. The fiducial power spectrum (PS) mean (black curves) and standard deviation (grey region), computed from 1000 realizations varying initial conditions and noise seeds at θ_{fid} . We note that the PS includes contributions from the cosmic signal, telescope noise, and wedge excision; cf. Fig. 1. The ‘data seed’ used as our mock observation is shown in red and is chosen to be a statistically ‘reasonable’ realization. The ‘model seed’ is shown in green and is chosen to be the realization closest to the mean, so that we can avoid varying the initial conditions when performing classical inference. The sharp drop seen at low k results from the wedge excision. All k -modes to the right of the vertical line $k \geq 0.2 \text{ Mpc}^{-1}$ are taken into account when computing the likelihood.




Figure 3. Correlation matrices of the 21-cm PS at the fiducial parameter choice, θ_{fid} (cf. equation 6). Upper-left: flattened 4D correlation matrix for the *mean removal* case with and without SKA noise (lower left and upper right triangles, respectively). Upper-right: zoom-in of the upper-left plot for the first three redshift bins. Bottom row: the same as the upper row, with addition of *+horizon cut*.

estimator to fit either the likelihood $\mathcal{L}(\mathbf{d}_{\text{PS}}|\theta)$ or the posterior $P(\theta|\mathbf{d}_{\text{PS}})$. We choose the former, which allows us to parametrize the form of the likelihood in some of our investigations, directly testing the impact of the previously mentioned approximations. Moreover, fitting the likelihood is more flexible as the priors can be changed afterwards without having to retrain. We can then perform inference using an explicit or non-explicit likelihood fitted with density estimation. This inference is then vastly more computationally efficient as the simulator (in our case 21CMFAST) is not called at each likelihood evaluation. This is why SBI is a form of so-called *amortized* inference: the computational cost is upfront in generating a training set for the density estimation. Once the likelihood is learned via density estimation, subsequent inference is

done very rapidly, with a single likelihood evaluation taking of order milliseconds.

Another and more popular form of amortized inference in 21-cm literature is emulation, where a deep neural network (NN) is trained to directly replace the simulator for pre-chosen summary statistics like the PS (e.g. Kern et al. 2017; Shimabukuro & Semelin 2017; Schmit & Pritchard 2018; Jennings et al. 2019; Ghara et al. 2020; Mondal et al. 2022; Breitman et al., in preparation). However, emulators themselves have errors that are difficult to include properly in the likelihood. With SBI, our only source of error (assuming the simulator is correct) is the density estimation, which can always be reduced by increasing the training set or using iterative resampling.




Figure 4. Redshift evolution of the 1D PS at $k = 0.25 \text{ Mpc}^{-1}$ (upper panel) and IGM mean neutral fraction (lower panel), covering the first two redshift bins. The red (blue) curves illustrate the mean computed over five realizations with large (small) power at $z = 6.7$. Consistent with Fig. 3, the panels demonstrate strong anticorrelation between $z = 5.9$ and $z = 6.7$, resulting from the relative differences in the steepness in the halo mass functions of these realizations, which impact the timing of the EoR (see text for details).

3.2.1 NDE loss function

Given an NN parametrization of a likelihood $\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta})$, one can train its parameters to approximate the real likelihood $\mathcal{L}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta})$. A natural distance between these distributions is the Kullback–Liebler (KL) divergence. For a continuous 1D random variable, with $p(x)$ as the true (target) distribution and $q(x)$ as the approximate distribution, the integral

$$D_{\text{KL}}(P\|Q) = \int_{-\infty}^{\infty} p(x) \log \frac{p(x)}{q(x)} dx \quad (11)$$

represents the information gain if the target distribution $p(x)$ were used instead of $q(x)$. Therefore, the main goal during training is to minimize this difference. The integral can be further written as

$$\begin{aligned} D_{\text{KL}}(P\|Q) &= \int_{-\infty}^{\infty} p(x) \log p(x) dx - \int_{-\infty}^{\infty} p(x) \log q(x) dx \\ &\approx -H_{\text{P}} - \frac{1}{N} \sum_{i=1}^N \log q(x_i), \end{aligned} \quad (12)$$

with samples x_i being drawn from $p(x)$. In the last expression, H_{P} is a constant denoting the Shannon entropy of the distribution $p(x)$ and the second term is the Monte Carlo estimation of the integral $\int_{-\infty}^{\infty} p(x) \log q(x) dx$.

We use equation (12) to define the loss between the NN likelihood estimate, $\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta}) \equiv q(x)$, and the true likelihood, $\mathcal{L}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta}) \equiv$

$p(x)$:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS},i}|\boldsymbol{\theta}_i), \quad (13)$$

where $\boldsymbol{\theta}_i$, $\mathbf{d}_{\text{PS},i}$ are samples from the training set. Once the parametric shape of \mathcal{L}_{NN} is defined (see the following section), one can minimize the loss via stochastic gradient descent (SGD) and back propagation.⁶

In this work, we pre-calculate a data base of $\approx 78\,000$ light-cones, pass them through the observational pipeline and compress to the power spectra (cf. Fig. 1). Realizations correspond to astrophysical parameters drawn from the prior distribution described in Section 2.1. Importantly, every realization has different cosmic and noise seeds. Around 57 000 are used for the training, 14 000 for validation, and 7000 for the final testing. We note that a data base of this size corresponds roughly to the number of samples (likelihood calls) needed for Markov chain Monte Carlo (MCMC) or nested inference to converge. We return to this point below, stressing that SBI is not only more accurate but also more computationally efficient compared to classical inference when the simulator is called on-the-fly (non-amortized).

3.2.2 NDE likelihood parametrizations

We use SBI to test the validity of the common approximations made when using a Gaussian likelihood. As discussed at the beginning of this section (Section 3.2), these include the following:

- (i) estimating the mean from a single realization;
- (ii) estimating the (co)variance only at the fiducial parameters;
- (iii) assuming a Gaussian form to begin with, thus ignoring higher order correlations.

Our aim is to investigate the consequences of each assumption, quantifying their impact on recovered posteriors. To do that, we relax constraints one by one, gradually enlarging the likelihood complexity and possibly improving final performance. For this, we use neural density estimators (NDEs) parametrized with simple fully connected layers.

In order to estimate the mean of the PS better, we use a feed-forward NN that takes parameters $\boldsymbol{\theta}$ and outputs the PS mean. In other words,

$$\boldsymbol{\mu}_{\text{NN}}(\boldsymbol{\theta}) = \text{NN}(\boldsymbol{\theta}). \quad (14)$$

The possible Gaussian likelihoods are then

$$\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta}) = \mathcal{N}(\mathbf{d}_{\text{PS}}|\boldsymbol{\mu}_{\text{NN}}(\boldsymbol{\theta}), \boldsymbol{\sigma}^2(\boldsymbol{\theta}_{\text{fid}})), \quad (15)$$

$$\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta}) = \mathcal{N}(\mathbf{d}_{\text{PS}}|\boldsymbol{\mu}_{\text{NN}}(\boldsymbol{\theta}), \boldsymbol{\Sigma}(\boldsymbol{\theta}_{\text{fid}})). \quad (16)$$

⁶In this work, we sample parameters from the prior $\pi(\boldsymbol{\theta})$ and then draw from the likelihood $\mathcal{L}(\mathbf{d}_{\text{PS}}|\boldsymbol{\theta})$ by simulating the PS. In this way, $\boldsymbol{\theta}_i$, $\mathbf{d}_{\text{PS},i}$ are sampled from the joint distribution $P(\boldsymbol{\theta}, \mathbf{d}_{\text{PS}})$. One could therefore think of defining the loss function as

$$L = -\frac{1}{N} \sum_{i=1}^N \pi(\boldsymbol{\theta}_i)^{-1} \log \mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS},i}|\boldsymbol{\theta}_i).$$

As we are interested in the distribution conditional on the parameter space (i.e. likelihood), both methods are correct. Introducing the factor $\pi(\boldsymbol{\theta}_i)^{-1}$ in the loss gives more weight to the samples from the prior edges, counterbalancing the fact that they are rare. We prefer to focus on the main prior volume, so we keep the loss as defined in the main text.

Table 1. Properties of the likelihoods used in this study. The CLASSIC likelihoods are used for non-amortized inference, in which the simulator is called at each likelihood evaluation. The NDE likelihoods make use of a pre-computed (amortized) set of simulations to fit explicit (first four) or non-explicit (final two) likelihoods. Based on the validation loss and SBC, we find *NDE Gauss mixture* to be the most accurate and stable to train.

	Non-Gaussian	Non-diagonal covariance	(Co)variance is a function of θ	Mean by averaging over realizations
CLASSIC fixed var	✗	✗	✗	✗(single, well-chosen seed)
CLASSIC fixed cov	✗	✓	✗	✗(single, well-chosen seed)
NDE fixed var	✗	✗	✗	✓
NDE fixed cov	✗	✓	✗	✓
NDE varying var	✗	✗	✓	✓
NDE varying cov	✗	✓	✓	✓
NDE CMAF	✓	N/A	N/A	N/A
NDE Gauss mixture	✓	N/A	N/A	N/A

Here $\sigma^2(\theta_{\text{fid}})$ and $\Sigma(\theta_{\text{fid}})$ represent the variance and covariance estimated at the fiducial parameter values, respectively. After training such likelihoods, the network will learn to interpolate PS mean values for the whole parameter space θ .

Likewise, we can also estimate the (co)variance matrix with an NN. In this scenario, the network can output one of the following:

$$\mu_{\text{NN}}(\theta), \sigma_{\text{NN}}^2(\theta) = \text{NN}(\theta), \quad (17)$$

$$\mu_{\text{NN}}(\theta), \Sigma_{\text{NN}}(\theta) = \text{NN}(\theta), \quad (18)$$

with their respective likelihoods:

$$\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\theta) = \mathcal{N}(\mathbf{d}_{\text{PS}}|\mu_{\text{NN}}(\theta), \sigma_{\text{NN}}^2(\theta)), \quad (19)$$

$$\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\theta) = \mathcal{N}(\mathbf{d}_{\text{PS}}|\mu_{\text{NN}}(\theta), \Sigma_{\text{NN}}(\theta)). \quad (20)$$

The last equation represents a likelihood complexity currently intractable by classical methods, as calculating a full covariance matrix at every point of the parameter space is too expensive.

Finally, we also adopt a non-explicit likelihood, allowing us to test the importance of the higher order terms that are ignored when assuming a Gaussian form. We try two non-explicit likelihood density estimators implemented in DELFI: (i) a Gaussian mixture network; and (ii) conditional masked autoregressive flow (CMAF). A Gaussian mixture density network is a simple density estimator that is very stable to train.⁷ For a mixture of K Gaussians, a feed-forward NN takes parameters θ and outputs

$$\mu_{\text{NN},1}(\theta), \Sigma_{\text{NN},1}(\theta), \dots, \mu_{\text{NN},K}(\theta), \Sigma_{\text{NN},K}(\theta), \Phi(\theta) = \text{NN}(\theta), \quad (21)$$

where $\mu_{\text{NN},i}(\theta)$, $\Sigma_{\text{NN},i}(\theta)$ describe mean and covariance of the i th Gaussian and $\phi_i(\theta) \in \Phi(\theta)$ its relative weight, where Φ is the vector of relative weights. Therefore, the full likelihood can be written as

$$\mathcal{L}_{\text{NN}}(\mathbf{d}_{\text{PS}}|\theta) = \sum_{i=1}^K \phi_i(\theta) \cdot \mathcal{N}(\mathbf{d}_{\text{PS}}|\mu_{\text{NN},i}(\theta), \Sigma_{\text{NN},i}(\theta)), \quad (22)$$

where $\sum_i \phi_i(\theta) = 1$. Such a parametrization will be able to catch both higher order moments and possible multimodalities in the likelihood. Here we use $K = 3$ G, as it provides the lowest validation loss.

We also use CMAF density estimation (Papamakarios, Pavlakou & Murray 2017; Papamakarios et al. 2019). A CMAF is parametrized

through a series of linear transformations of normal random variables. The idea is to start from a unit variance, zero-mean normal distribution $p_u(\mathbf{u})$ and through the transform $\mathbf{x} = T(\mathbf{u})$ find the estimate of some final distribution $p_x(\mathbf{x})$. If the transformation T is invertible and differentiable, one can write the final distribution as

$$\begin{aligned} p_x(\mathbf{x}) &= p_u(\mathbf{u}) |\det J_T(\mathbf{u})|^{-1} \\ &= p_u(T^{-1}(\mathbf{u})) |\det J_{T^{-1}}(\mathbf{x})|, \end{aligned} \quad (23)$$

where J_T and $J_{T^{-1}}$ are Jacobians of T and its inverse, respectively. For larger expressivity, one can stack several transformations so that $T = T_n \circ T_{n-1} \circ \dots \circ T_1$. This will reflect in the above equation as the multiplication of n Jacobian determinants. The parameters of the network are then parameters of the transformations, trained in the same way as before – by minimizing KL divergence between transformed distribution and the real one. In our case, $\mathbf{x} = (\mathbf{d}_{\text{PS}}, \theta)$ represents the space of the PS and parameters. In order to construct a likelihood, one then fixes parameter dimensions and transforms PS dimensions only. For more details, see a recent review by Papamakarios et al. (2021). Although CMAF NDEs offer more flexibility to capture complicated likelihoods, they are less stable to train compared to Gaussian mixture NDEs. Indeed, we use Gaussian mixture as our reference NDE below, as it has the best validation loss (see results) and is the least biased (see Section 4.2.2).

For all NDEs except CMAF, the underlying NNs are defined with 10 fully connected layers of 100 neurons, LEAKYRELU (e.g. Xu et al. 2015) activation function (except for the last layer) followed by a layer defining the likelihood parameters. The CMAF NDE is of the same shape, using the same activation function, where the sequence of layers parametrizes the chain of transformations. Our NNs use the Adam optimizer (Kingma & Ba 2014), learning rate reduction, and early stopping. All NDE likelihoods and training procedures are made public as a part of 21CMLIKELIHOODS⁸ package.

Finally, it is important to test that the likelihood density estimation has trained well, without introducing biases and being under/overconfident. One can use a mock observation and compare the resulting posteriors obtained from different likelihood estimators. If the mock observation is not consistent with the recovered posterior, then the likelihood estimator is biased (Section 4.1). If this is repeated for many mock observations, sampled from the joint distribution $P(\theta, \mathbf{d}_{\text{PS}})$, the average over all posteriors should reproduce the prior. This method (described in more details in Section 4.2.2) is called simulation-based calibration (SBC; Talts et al. 2018). SBC provides a rigorous, quantitative test of the accuracy of likelihood estimation.

⁷A Gaussian mixture density network should not be confused with a Gaussian mixture model. The latter describes a distribution where each Gaussian represents samples of a particular class. This is not the case here – the mixture density network is used purely for better parametrization of the likelihood.

⁸<https://github.com/dprelogo/21cmLikelihoods>




Figure 5. Posterior of the fiducial mock observation recovered from the *NDE Gauss mixture* likelihood, compared with *CLASSIC fixed cov* and *CLASSIC fixed var*. Individual figures show 2D or 1D marginalizations of the posterior. Contours label 95 per cent confidence interval (CI), and black-dashed lines true parameter values.




Figure 6. Like Fig. 5, but here we compare the posteriors from *NDE fixed cov* and *NDE fixed var*.

Unfortunately, it requires averaging over thousands of posteriors and is therefore only computationally tractable for amortized inference like SBI.

4 RESULTS

In summary, we compare the performance of the following likelihood choices for the spherically averaged 21-cm PS.

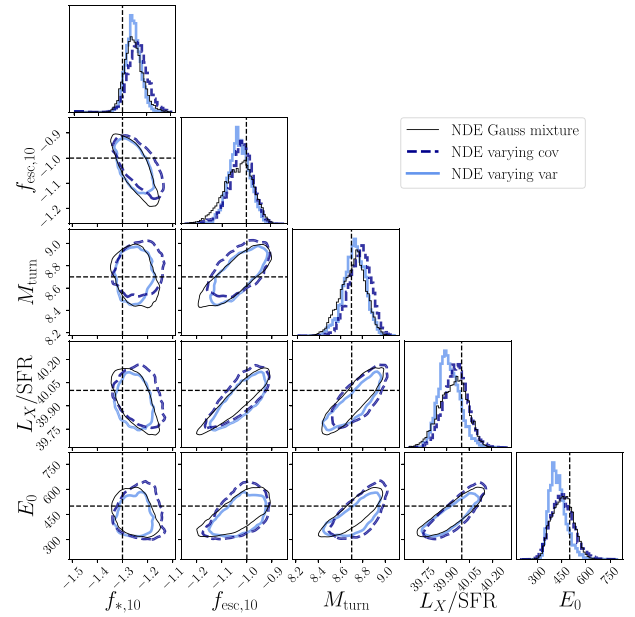


Figure 7. Like Fig. 5, but here we compare the posteriors from *NDE varying cov* and *NDE varying var*.




Figure 8. Like Fig. 5, but here comparing the two non-Gaussian likelihoods, *NDE Gauss mixture* and *NDE CMAF*.

(i) *CLASSIC fixed var* – assumes a Gaussian likelihood, with the mean estimated from a single (albeit well chosen) realization, and a diagonal covariance estimated only at θ_{fid} . This is by far the most common choice in 21-cm inference. In fact, this is an improvement over the most common method of *randomly* picking a realization for the model. As mentioned above, we pick the one realization out of 1000 that is closest to the mean at the fiducial parameter choice.

(ii) *CLASSIC fixed cov* – like *CLASSIC fixed var*, but using a non-diagonal covariance.




Figure 9. Summary of the previous posteriors, showing only 1D marginal PDFs for all likelihood choices. The fiducial (true) parameter values are denoted with horizontal, dashed lines. Most posteriors are overconfident. The most accurate posterior estimate (see Section 4.2.2) comes from the *NDE Gauss mixture*.

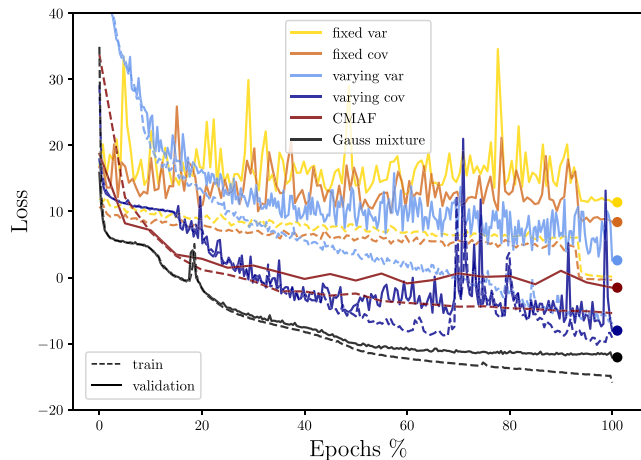


Figure 10. Training and validation losses for all NDEs. Loss on the y-axis corresponds to an unnormalized KL divergence and can be used as a relative measure of accuracy (see text for details). On the x-axis we show the per cent of the total number of epochs used for training (100 per cent corresponds to the minimum of the validation loss denoted with dots).

(iii) *NDE fixed var* – assumes a Gaussian likelihood, with the mean estimated by NDE, and a diagonal covariance estimated only at θ_{fid} .

(iv) *NDE fixed cov* – like *NDE fixed var*, but with a non-diagonal covariance.

(v) *NDE varying var* – assumes a Gaussian likelihood, with the mean and a diagonal covariance estimated throughout parameter space by NDE.

(vi) *NDE varying cov* – like *NDE varying var*, but with the NDE fitting a non-diagonal covariance.

(vii) *NDE CMAF* – a non-Gaussian likelihood estimator using CMAF NDE.

(viii) *NDE Gauss mixture* – a non-Gaussian likelihood estimator using Gaussian mixture NDE.

For the first two ‘classic’ likelihood choices, the simulator is called on-the-fly during inference (i.e. it is not amortized). The other likelihoods are used in SBI (amortized inference). Using validation loss and SBC rank statistics, (see Figs 10 and 12), we determine that the *NDE Gauss mixture* is the most accurate of the likelihood estimators considered here. We therefore use it as a reference throughout the rest of this work. We summarize the assumptions for each likelihood choice in Table 1.

We compare the performances of the likelihood choices in two parts. In Section 4.1, we show posteriors for the mock observation, and qualitatively describe their performances. In Section 4.2, we quantitatively test all NDE likelihoods using many mock observations across parameter space.

4.1 Posteriors for the fiducial mock observation

In order to recover the posteriors of the fiducial mock observation (see Fig. 2), we run Bayesian inference for the constructed and trained likelihoods. For the two classic choices, we use 21CMMC⁹ package, together with the MULTINEST¹⁰ sampler (Buchner et al. 2014). For the

⁹<https://github.com/21cmfast/21CMMC>

¹⁰<https://github.com/JohannesBuchner/PyMultiNest>

NDEs, we use ULTRANEST¹¹ (Buchner 2021), with the MLFriends algorithm (Buchner 2016, 2019). For simple posteriors like ours the two codes perform the same; however, the computational cost of SBI is lower with ULTRANEST’s vectorized likelihood evaluation. For our default convergence criteria that require $\sim 40\,000$ – $90\,000$ likelihood evaluations, the non-amortized (classic) inference with 21CMMC takes $\sim 10^5$ core hours, while the amortized (SBI) inference with ULTRANEST takes ~ 0.1 core hours. We return to the relative computational costs in more detail below.

In Fig. 5, we show 1D and 2D marginal posteriors for the two classic likelihoods in green, together with *NDE Gauss mixture* as a reference in black. Straight lines denote the fiducial parameter values and 95 per cent confidence interval (CI) are demarcated in the 2D panels. Comparing *CLASSIC fixed var* with *CLASSIC fixed cov*, we can clearly see that ignoring the non-diagonal covariance terms results in a higher bias and a more overconfident posterior. This is in contrast with Zhao et al. (2022b) who find that for their two parameter model, the NDE likelihood results in a *tighter* posterior compared with a Gaussian likelihood. This difference could be due to the specific choice of model seed and/or the calculation of the variance (computed via Monte Carlo versus assumed to be Poisson). Regardless, we agree with their main conclusions that just using a diagonal Gaussian variance can bias inference. Our true parameter values are in some cases inconsistent with the *CLASSIC fixed var* posterior, *the most popular choice for 21-cm inference*, at greater than 95 per cent CI. Note that we confirm that our choice of data seed is not a rare outlier (see Fig. 2), and so the truth should be inside the 68 per cent CI. This bias is especially obvious for the X-ray heating parameters, E_0 and L_X/SFR . Including a full covariance, still only evaluated at θ_{fid} , gives a more accurate recovery. Indeed, the true values of the fiducial mock are always consistent with the recovered *CLASSIC fixed cov* posterior. Because it only requires pre-computing a few hundred additional realizations of the fiducial model, the *CLASSIC fixed cov* likelihood provides an easy way of increasing the precision of classical inference. However, even though it is a significant improvement over *CLASSIC fixed var*, we see that *CLASSIC fixed cov* is still too narrow compared with our reference, non-Gaussian posterior in black.

In Fig. 6, we show results for *NDE fixed cov* in brown and *NDE fixed var* in yellow, with *NDE Gauss mixture* for reference as before. Comparing to the classic likelihood results from Fig. 5, we can see that the bias and constraints from *NDE fixed var* are consistent with *CLASSIC fixed var*, while those from *NDE fixed cov* are consistent with *NDE fixed var*. There is a difference in the direction of the bias, which is expected as the mean in the two cases is estimated by different methods. We therefore conclude that using NDE to find the mean PS does not provide a significant improvement over using a single realization as the mean, *provided that the realization is carefully chosen to be the one closest to the mean at the fiducial parameter value*. This last caveat is important, as most classic inferences *randomly* choose an initial condition seed to be used in inference.

In Fig. 7, we show posteriors from *NDE varying var* and *NDE varying cov*, which allow the (co)variance to be a function of θ . In comparison to the previous results for which the (co)variance is only evaluated at a single point, we see a notable reduction in bias. In fact, *NDE varying cov* results in a very similar posterior as our reference *NDE Gauss mixture*. This suggests that higher order moments (beyond the two-point covariance) do not contribute

¹¹<https://github.com/JohannesBuchner/UltraNest>




Figure 11. Training and validation losses for the *NDE Gauss mixture* model, for different training set sizes (denoted as a percentage of the full set, 57 000, in the legend). As in Fig. 10, the loss on the y-axis is the unnormalized KL divergence and the x-axis shows the training percentage.

significantly to the 21-cm PS likelihood. As we demonstrate in the next section, *NDE varying cov* also generalizes well across our parameter space of mock observations, making it our second most accurate likelihood.

Finally, in Fig. 8, we compare the posteriors from our two non-Gaussian likelihoods: *NDE Gauss mixture* and *NDE CMAF*. Surprisingly, our CMAF results in an overconfident and biased posterior. The fact that the CMAF is more expressive (i.e. is able to adapt to many different distribution shapes) makes it more difficult to train, requiring either a larger training set or sequential (active) training (Papamakarios et al. 2019; Alsing et al. 2019). Thankfully, the 21-cm PS likelihood is simple enough to be well fit by the more stable *NDE Gauss mixture*, and does not require the additional flexibility of CMAF density estimation.

We summarize our results in Fig. 9, by showing 1D distributions for all parameters (rows) and all likelihoods (columns). Black dashed lines denote the true parameter values used for the mock observation. For the *NDE Gauss mixture* the true values are consistent with all of the 1D marginal PDFs. Although we cannot quantitatively judge the likelihood using only one posterior, we show in the following section that *NDE Gauss mixture* performs the best in general. *NDE varying cov* follows it very closely, in both confidence and bias. *NDE varying var* performs slightly worse. The largest biases are seen for the fixed variance likelihoods. These results confirm that it is important to account for the full covariance between wave modes and redshifts, preferably at multiple points in parameter space. Higher order moments are not particularly important for characterizing the 21-cm PS likelihood.

4.2 Do the likelihood estimators generalize well across parameter space?

In the previous section, we compared the posteriors resulting from our different likelihood estimators for a single mock observation. In order to gain confidence in the ability of the likelihoods to generalize across parameter space, we must compare the posteriors for many mock observations generated at many different θ . Here we quantify this using two metrics: (i) the validation loss; and (ii) rank distributions from SBC. We discuss each in turn.

Unfortunately, it is impractical to generate thousands of classical, non-amortized inferences required for this assessment. We therefore only use our SBI likelihoods in this section. As was demonstrated in the previous section, *NDE fixed var* and *NDE fixed cov* can serve as proxies for *CLASSIC fixed var* and *CLASSIC fixed cov*, respectively.

4.2.1 Training and validation loss

Our loss function (equation 13) is an unnormalized estimate of the KL divergence between likelihood samples and the fitted distribution. For this reason, the validation loss is a measure of how well the fitted likelihood function generalizes over the whole parameter space. Because of the missing normalization constant (Shannon entropy of the ‘real’ distribution, see equation 12), the loss function can be negative, while the KL divergence cannot.

Fig. 10 shows the training and validation losses for all NDEs considered in this work, in dashed and solid lines, respectively. On the x-axis we show the percentage of the total number of epochs used for the final training (100 per cent corresponds to the minimum of the validation loss). The relative validation losses follow the same trends as noted in the previous section for the case of a single mock observation. Namely, *NDE Gauss mixture* has the lowest validation loss, confirming that it can accurately fit the likelihood across our parameter space. *NDE varying cov* is the second best, while *NDE fixed var* is the worst.

It is useful to know just how large of a training set is needed to achieve the accuracy demonstrated by *NDE Gauss mixture*. Our default training set has 57 000 simulations. This is roughly comparable to the average number of likelihood evaluations needed for ULTRANEST or MULTINEST to converge when doing inference.¹² The fact that these numbers are comparable means that amortized (running simulations before inference) and non-amortized (running simulations during inference) approaches to inference have comparable computational costs. However, if our best performing *NDE Gauss mixture* estimator can be trained with a smaller training set, then it would not only be more accurate but also faster than classical, non-amortized inference.

To answer this question, we retrained *NDE Gauss mixture* using only a subset of our training set: 50 per cent (28 000 simulations) and 10 per cent (5700 simulations).¹³ The training and validation losses are plotted in Fig. 11. One can clearly see that reducing the training size by a factor of produces little change in the validation loss. Therefore, 20 000–30 000 simulations are sufficient for *NDE Gauss mixture* to maintain accuracy. Reducing the training size by a factor of 10 pushes the validation loss to roughly the level of *NDE varying cov* (see Fig. 10). Since this is our second best model, good performance is obtained even with ~ 6000 simulations. We thus conclude that SBI

¹²The exact number of likelihood evaluations depends not only on the mock observation but also on the sampler settings such as the number of live points and the convergence criterion. Our default setting uses 400 live points, requiring $\sim 90\,000$ likelihood evaluations. We find that using 200 live points can give similar posteriors for our fiducial mock, requiring $\sim 40\,000$ likelihood evaluations. Using a smaller number of live points results in noisy/biased posteriors.

¹³For visualization purposes, we keep the validation set fixed (17 000 samples). This insures a smooth evolution of the validation loss curve; however, we confirm that such a large validation set is not needed. We tested that for our usage case, having 1000 samples was more than enough for the simple purpose of detecting overfitting and reducing the learning rate if needed.

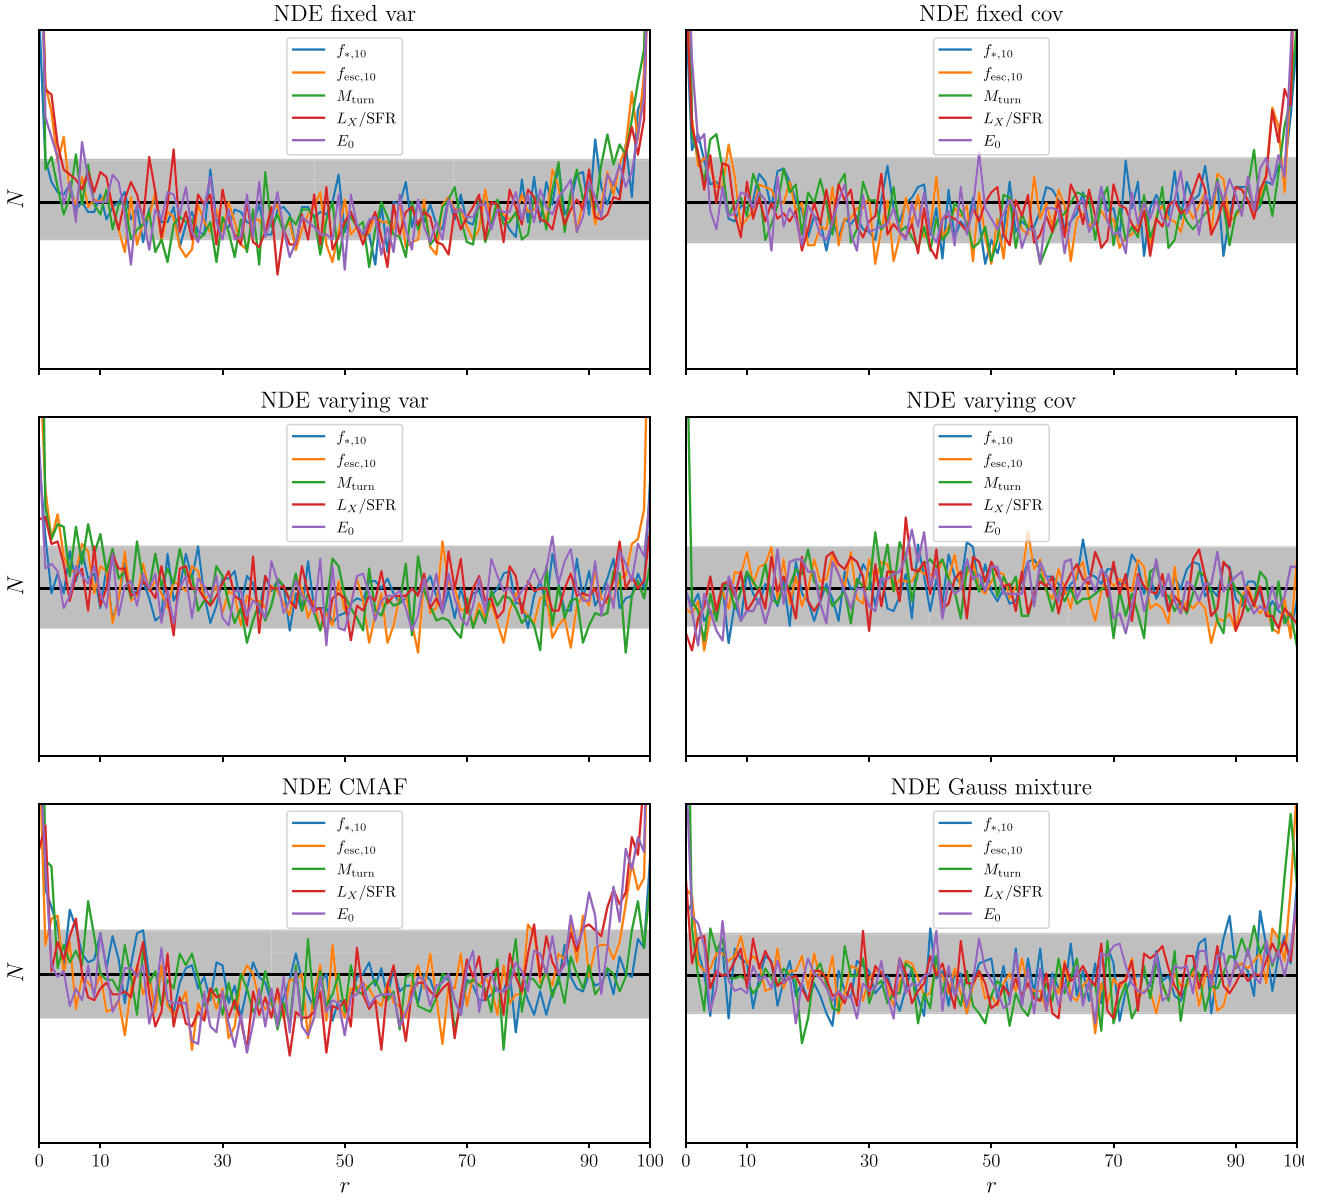


Figure 12. Posterior rank distributions from SBC computed from ~ 7000 models in the test set. Each panel corresponds to a different NDE likelihood. The x-axis denotes the rank r from 0 to $L = 100$. The y-axis shows the posterior number count N of a particular rank. Black line and shaded area mark the mean and 95 per cent confidence interval (CI) of the binomial distribution, respectively.

is not only more accurate than classic, non-amortized inference, but also is faster by factors of $\gtrsim 3-10$.

The validation loss discussed in this section quantify the quality of the likelihood fits, not the performance of the whole Bayesian inference framework. In the following section, we test the accuracy of the recovered posteriors and thus the robustness of our likelihood functions.

4.2.2 Simulation-based calibration

In Section 4.1, we showed posteriors obtained from a single mock observation. We compared the performances of different likelihood choices by noting if the true values from the mock were consistent with the recovered posteriors. Here we extend this comparison

over many different mock observations, in a procedure known as simulation-based calibration (SBC).

The core of the method is based on the following steps:

- (i) sample parameters from prior, $\tilde{\theta} \sim \pi(\theta)$;
- (ii) sample the likelihood with a simulator, $\tilde{\mathbf{d}} \sim \mathcal{L}(\mathbf{d}|\tilde{\theta})$;
- (iii) calculate the posterior using the sample as a mock observation, $\tilde{\mathbf{d}} \rightarrow P(\theta|\tilde{\mathbf{d}})$.

Here the samples are labelled with \sim . The first two steps are identical to SBI, with the final step adding the calculation of the posterior corresponding to that mock observation. If we repeat these steps many times and average the posteriors for many $(\tilde{\theta}, \tilde{\mathbf{d}})$, it

follows that

$$\begin{aligned}\pi(\boldsymbol{\theta}) &= \int P(\boldsymbol{\theta}|\tilde{\mathbf{d}}) \cdot \mathcal{L}(\tilde{\mathbf{d}}|\tilde{\boldsymbol{\theta}})\pi(\tilde{\boldsymbol{\theta}}) d\tilde{\boldsymbol{\theta}} d\tilde{\mathbf{d}} \\ &\approx \frac{1}{N} \sum_{i=1}^n P(\boldsymbol{\theta}|\tilde{\mathbf{d}}_i),\end{aligned}\quad (24)$$

where the last equation is the Monte Carlo estimate of the integral for samples drawn from $P(\boldsymbol{\theta}, \mathbf{d}) = \mathcal{L}(\mathbf{d}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$. In other words, *prior = data-averaged posterior*. If any step of the Bayesian inference pipeline went wrong, the equality will not hold. This is the basis of SBC (Talts et al. 2018). In particular, if we sample L points $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_L)$ from the posterior $P(\boldsymbol{\theta}|\tilde{\mathbf{d}})$, make a scalar function $f : \Theta \rightarrow \mathbb{R}$, and define a rank function,

$$r(f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_L), f(\tilde{\boldsymbol{\theta}})) = \sum_{i=1}^L \mathbb{1}[f(\boldsymbol{\theta}_i) < f(\tilde{\boldsymbol{\theta}})], \quad (25)$$

a distribution of a rank r should be uniform in $[0, L]$. Here the Boolean $\mathbb{1}$ equals 1 if the condition is true, and 0 otherwise. Therefore, the *rank* of a particular posterior counts how many samples are below its true value. The *rank distribution* for rank r counts how many posteriors have r samples out of L below their true values. This equation is a 1D representation of the data-averaged posterior in equation (24). By testing the uniformity in the rank r , one can see if posteriors are statistically underconfident (overconfident) or biased. Overconfidence will make distribution convex (\cup -shaped), underconfidence will make distribution concave (\cap -shaped), and any bias will tilt the distribution to one side. For instance, having a \cap -shaped profile in the case of underconfident posteriors can be intuitively understood as the following. Rank statistics measures the difference between the prior CDF and the data-averaged posterior CDF. If the two distributions are the same, the rank statistics are zero and flat. However if the data-averaged posterior is wider, the low-rank points would move towards the middle ranks, as it is harder to obtain low ranks for a wide distribution. The opposite is true for the overconfident posteriors. For more details, see Talts et al. (2018).

We consider five scalar functions $f_i(\boldsymbol{\theta}) = \theta_i$, one for each astrophysical parameter. This allows us to assess the quality of the 1D posteriors for each parameter separately. In order to calculate the rank statistics, we run ULTRANEST for all likelihoods and all samples in our test set of ~ 7000 simulations. ULTRANEST is the best choice for this task because of its (i) vectorization of the likelihood calculation; (ii) robust convergence criteria; and (iii) non-correlated samples in the final posterior. The last point is crucial for SBC as correlated samples contaminate the rank statistics. We calculate the ranks of the posterior samples via equation (25), for $L = 100$.

The results are shown in Fig. 12. Each rank statistic is expected to follow a binomial distribution – in the plot, the black line and grey area show its mean and 95 per cent interval, respectively. In general, the quality of the 1D posteriors is comparable within each likelihood choice (there is no obvious outlying astrophysical parameter). For *NDE fixed var* and *NDE fixed cov* we can see a convex distribution, indicating that their posteriors are overconfident. In the case of *NDE varying var* and *NDE varying cov*, results look much better, with the curves staying in the 95 per cent region almost throughout the rank statistics. Furthermore, *NDE CMAF* seems to be fairly overconfident and slightly biased to the right, especially for recovery of L_X/SFR and E_0 .

NDE Gauss mixture performs the best across most of the rank statistic range, followed by *NDE varying cov*. The former seems slightly overconfident while the latter slightly underconfident, but the distributions are generally consistent with the expected binomial

95 per cent interval in grey. This confirms the previous results that *NDE Gauss mixture* provides the most accurate likelihood.

5 CONCLUSIONS

The spherically averaged PS is the most used summary statistic when performing inference from 21-cm tomography (using either current mock or upcoming data). However, the likelihood functions commonly used in such inferences contain several questionable assumptions. Almost exclusively a Gaussian functional form is adopted, despite the fact the signal is non-Gaussian. Even after assuming a Gaussian likelihood, further simplifications are made to make the likelihood evaluation computationally tractable. These include: (i) estimating the mean from a single, random realization; (ii) assuming a diagonal covariance; and (iii) computing the (co)variance only at a single point in parameter space.

Here, we systematically test these assumptions with SBI: we train NDEs to fit the likelihood using a data base of tens of thousands of simulated 21-cm PS observations. Our simulation pipeline consists of the following steps (cf. Fig. 1): (i) sample the initial matter PS and five galaxy parameters, computing a corresponding realization of the cosmic 21-cm light-cone with 21CMFAST; (ii) add a realization of the telescope noise corresponding to a 1000 h integration with SKA1-Low; (iii) zero all Fourier modes lying within a foreground-dominated horizon wedge; and (iv) bin the light-cone in redshift intervals, computing the 1D 21-cm PS for each bin. We make a data base of 57 000 such realizations of the PS for training, including another 14 000 for validation and 7000 for testing. We compare the results of inference for various likelihood choices, validating the ability of our NDEs to generalize over parameter space using SBC.

Our main results are the following.

(i) The most common likelihood choice – a Gaussian with a fixed variance evaluated only at a fiducial parameter – results in an overconfident and biased posterior (see also Zhao et al. 2022b). True values can be outside of the recovered 95 per cent CI when they should be inside 68 per cent CI.

(ii) Including the two-point covariance between Fourier modes and redshifts dramatically improves inference. If the covariance is fit as a function of the parameters, the resulting posteriors are almost identical to our best-performing non-Gaussian NDE. However, even pre-computing the covariance only at the fiducial parameter set gives decent results, making it a computationally viable option for classic, non-amortized inference.

(iii) A single realization can be used in place of the mean PS across parameter space, provided that the realization is pre-selected to lie close to the mean at the fiducial parameter value.

(iv) Our most accurate posteriors result from a non-Gaussian likelihood, using a simple-to-train Gaussian mixture NDE. More complex non-Gaussian NDEs like CMAFs are unnecessary and more difficult to train.

(v) Our best-performing likelihood estimator can be trained on a data base of only $\sim 20\,000$ – $30\,000$ simulations without any loss of accuracy. Good performance (comparable to our second best estimator) can be achieved with only 6000 simulations. This is up to an order of magnitude smaller than the number of likelihood evaluations needed for inference. Therefore creating a training set to fit the likelihood *before* running inference is more computationally efficient than calling a simulator *on-the-fly* with each likelihood evaluation during inference.

Although the choice of likelihood is unlikely to be important for current noisy 21-cm observations, we should have a high S/N

detection of the cosmic 21-cm PS in the near future. This work highlights the power of SBI to provide accurate posteriors at a comparably low computational cost.

ACKNOWLEDGEMENTS

We thank B. Greig, A. Liu, S. Murray, Y. Qin, and B. Wandelt for helpful comments on a draft version of this paper. AM acknowledges support from the Ministry of Universities and Research (MUR) through the PNRR project ‘Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing’ and the PRO3 project ‘Data Science methods for Multi-Messenger Astrophysics and Cosmology’. We gratefully acknowledge computational resources of the HPC center at SNS. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center (PSC).

DATA AVAILABILITY

The data underlying this paper will be shared on reasonable request to the corresponding author.

REFERENCES

- Abdurashidova Z. et al., 2022a, *ApJ*, 924, 51
 Abdurashidova Z. et al., 2022b, *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
 Barkana R., Loeb A., 2004, *ApJ*, 609, 474
 Barkana R., Loeb A., 2005, *ApJ*, 626, 1
 Bayer A. E., Modi C., Ferraro S., 2023, *J. Cosmol. Astropart. Phys.*, 06, 046
 Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
 Bouwens R. J. et al., 2015a, *ApJ*, 803, 34
 Bouwens R. J., Illingworth G. D., Oesch P. A., Caruana J., Holwerda B., Smit R., Wilkins S., 2015b, *ApJ*, 811, 140
 Bouwens R. J., Oesch P. A., Illingworth G. D., Ellis R. S., Stefanon M., 2017, *ApJ*, 843, 129
 Buchner J. et al., 2014, *A&A*, 564, A125
 Buchner J., 2016, *Stat. Comput.*, 26, 383
 Buchner J., 2019, *PASP*, 131, 108005
 Buchner J., 2021, *J. Open Source Softw.*, 6, 3001
 Cole A., Miller B. K., Witte S. J., Cai M. X., Grootes M. W., Nattino F., Weniger C., 2022, *J. Cosmol. Astropart. Phys.*, 09, 004
 Cranmer K., Brehmer J., Louppe G., 2020, *Proc. Natl. Acad. Sci.*, 117, 30055
 Dai B., Seljak U., 2022, *MNRAS*, 516, 2363
 Das A., Mesinger A., Pallottini A., Ferrara A., Wise J. H., 2017, *MNRAS*, 469, 1166
 Fragos T. et al., 2013, *ApJ*, 764, 41
 Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, *ApJ*, 613, 1
 Gazagnes S., Koopmans L. V. E., Wilkinson M. H. F., 2021, *MNRAS*, 502, 1816
 Ghara R. et al., 2020, *MNRAS*, 493, 4728
 Giri S. K., Schneider A., Maion F., Angulo R. E., 2023, *A&A*, 669, A6
 Giri S., Mellema G., Jensen H., 2020, *J. Open Source Softw.*, 5, 2363
 Greig B. et al., 2021, *MNRAS*, 501, 1
 Greig B., Mesinger A., 2015, *MNRAS*, 449, 4246
 Greig B., Mesinger A., 2018, *MNRAS*, 477, 3217
 Greig B., Ting Y.-S., Kurov A. A., 2022, *MNRAS*, 513, 1719
 Greig B., Ting Y.-S., Kurov A. A., 2023, *MNRAS*, 519, 5288
 HERA Collaboration et al., 2023, *ApJ*, 945, 124
 Jasche J., Kitaura F. S., 2010, *MNRAS*, 407, 29
 Jasche J., Wandelt B. D., 2012, *MNRAS*, 425, 1042
 Jasche J., Wandelt B. D., 2013, *MNRAS*, 432, 894
 Jennings W. D., Watkinson C. A., Abdalla F. B., McEwen J. D., 2019, *MNRAS*, 483, 2907
 Jensen H. et al., 2013, *MNRAS*, 435, 460
 Kern N. S., Liu A., Parsons A. R., Mesinger A., Greig B., 2017, *ApJ*, 848, 23
 Kingma D. P., Ba J., 2014, preprint (arXiv:1412.6980)
 Kitaura F. S., EnBlin T. A., 2008, *MNRAS*, 389, 497
 Koopmans L. et al., 2015, PoS, AASKA14, 001
 Leclercq F., Jasche J., Lavaux G., Wandelt B., Percival W., 2017, *J. Cosmol. Astropart. Phys.*, 06, 049
 Lehmer B. D. et al., 2016, *ApJ*, 825, 7
 Liu A., Parsons A. R., Trott C. M., 2014a, *Phys. Rev. D*, 90, 023018
 Liu A., Parsons A. R., Trott C. M., 2014b, *Phys. Rev. D*, 90, 023019
 Liu A., Shaw J. R., 2020, *PASP*, 132, 062001
 Loeb A., Zaldarriaga M., 2004, *Phys. Rev. Lett.*, 92, 211301
 Maity B., Choudhury T. R., 2023, *MNRAS*, 521, 4140
 Mao Y., Shapiro P. R., Mellema G., Iliev I. T., Koda J., Ahn K., 2012, *MNRAS*, 422, 926
 McAlpine S. et al., 2022, *MNRAS*, 512, 5823
 McGreer I. D., Mesinger A., D’Odorico V., 2015, *MNRAS*, 447, 499
 McQuinn M., D’Aloisio A., 2018, *J. Cosmol. Astropart. Phys.*, 10, 016
 McQuinn M., O’Leary R. M., 2012, *ApJ*, 760, 3
 Mellema G. et al., 2013, *Exp. Astron.*, 36, 235
 Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
 Mesinger A., 2020, *The Cosmic 21-cm Revolution: Charting the First Billion Years of Our Universe*. IoP Publishing, Bristol
 Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663
 Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
 Mineo S., Gilfanov M., Sunyaev R., 2012, *MNRAS*, 419, 2095
 Mirocha J., Furlanetto S. R., 2019, *MNRAS*, 483, 1980
 Mondal R., Bharadwaj S., Majumdar S., 2017, *MNRAS*, 464, 2992
 Mondal R., Mellema G., Murray S. G., Greig B., 2022, *MNRAS*, 514, L31
 Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, *ApJ*, 752, 137
 Muñoz J. B., 2023, *MNRAS*, 523, 2587
 Murray S. G., Trott C. M., 2018, *ApJ*, 869, 25
 Murray S., Greig B., Mesinger A., Muñoz J., Qin Y., Park J., Watkinson C., 2020, *J. Open Source Softw.*, 5, 2582
 Nasirudin A., Murray S. G., Trott C. M., Greig B., Joseph R. C., Power C., 2020, *ApJ*, 893, 118
 Nikolić I., Mesinger A., Qin Y., Gorce A., 2023, preprint (arXiv:2307.01265)
 O’Shea B. W., Wise J. H., Xu H., Norman M. L., 2015, *ApJ*, 807, L12
 Oesch P. A., Bouwens R. J., Illingworth G. D., Labbé I., Stefanon M., 2018, *ApJ*, 855, 105
 Papamakarios G., Nalisnick E., Jimenez Rezende D., Mohamed S., Lakshminarayanan B., 2021, *The Journal of Machine Learning Research*, 22, 2617
 Papamakarios G., Pavlakou T., Murray I., 2017, *Advances in neural information processing systems*, 30
 Papamakarios G., Sterratt D. C., Murray I., 2019, *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, p. 837
 Park J., Mesinger A., Greig B., Gillet N., 2019, *MNRAS*, 484, 933
 Parsons A. R. et al., 2014, *ApJ*, 788, 106
 Planck Collaboration VI, 2020, *A&A*, 641, A6
 Planck Collaboration XLVII, 2016, *A&A*, 596, A108
 Pober J. C. et al., 2014, *ApJ*, 782, 66
 Prelogović D., Mesinger A., Murray S., Fiameni G., Gillet N., 2022, *MNRAS*, 509, 3852
 Pritchard J. R., Furlanetto S. R., 2007, *MNRAS*, 376, 1680
 Saxena A., Cole A., Gazagnes S., Meerburg P. D., Weniger C., Witte S. J., 2023, preprint (arXiv:2303.07339)
 Schmit C. J., Pritchard J. R., 2018, *MNRAS*, 475, 1213
 Schneider A., Giri S. K., Mirocha J., 2021, *Phys. Rev. D*, 103, 083025
 Scoccimarro R., 1998, *MNRAS*, 299, 1097
 Shaw A. K., Bharadwaj S., Mondal R., 2019, *MNRAS*, 487, 4951
 Shimabukuro H., Semelin B., 2017, *MNRAS*, 468, 3869
 Shimabukuro H., Yoshiura S., Takahashi K., Yokoyama S., Ichiki K., 2017, *MNRAS*, 468, 1542

- Sobacchi E., Mesinger A., 2013a, *MNRAS*, 432, L51
Sobacchi E., Mesinger A., 2013b, *MNRAS*, 432, 3340
Sobacchi E., Mesinger A., 2014, *MNRAS*, 440, 1662
Talts S., Betancourt M., Simpson D., Vehtari A., Gelman A., 2018, preprint (arXiv:1804.06788)
Trott C. M. et al., 2020, *MNRAS*, 493, 4711
Trott C. M., Wayth R. B., Tingay S. J., 2012, *ApJ*, 757, 101
Vedantham H., Udaya Shankar N., Subrahmanyam R., 2012, *ApJ*, 745, 176
Watkinson C. A., Greig B., Mesinger A., 2022, *MNRAS*, 510, 3838
- Xu B., Wang N., Chen T., Li M., 2015, preprint (arXiv:1505.00853)
Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, *ApJ*, 833, 84
Zhao X., Mao Y., Cheng C., Wandelt B. D., 2022a, *ApJ*, 926, 151
Zhao X., Mao Y., Wandelt B. D., 2022b, *ApJ*, 933, 236

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.