

# Generative models of 21 cm Epoch of Reionisation lightcones with 3D scattering transforms

Ian Hothi<sup>1,2,\*</sup>, Erwan Allys<sup>1</sup>, Benoit Semelin<sup>2</sup>, and Romain Meriot<sup>3</sup>

<sup>1</sup> Laboratoire de Physique de l'ENS, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, 75005 Paris, France

<sup>2</sup> LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Université, F-75014 Paris, France

<sup>3</sup> Department of Physics, Blackett Laboratory, Imperial College London, London SW7 2AZ, UK

Received 11 June 2025 / Accepted 26 November 2025

## ABSTRACT

The 21 cm signal from the Epoch of Reionisation (EoR) is observed as a 3D dataset known as a lightcone, consisting of two spatial sky plane axes and a redshift (frequency) axis. Owing to its strongly non-Gaussian nature, fully characterising this signal requires summary statistics that go beyond two-point power spectra statistics. Recent developments in astrophysics, particularly in the context of the Galactic interstellar medium, demonstrate the efficacy of scattering transforms—novel summary statistics—to characterise fields with highly non-Gaussian properties. In particular, these statistics allow us to construct maximum-entropy generative models, even from a single target map, from which we can sample new, almost statistically identical realisations of a given process. Motivated by these advances, we extended the scattering transform formalism from 2D datasets to 3D EoR lightcones. To this end, we introduced a 3D wavelet set from the tensor product of 2D isotropic wavelets in the sky plane domain and 1D wavelets in the redshift domain. To test how well this 3D scattering transform can characterise an EoR lightcone, we constructed a maximum entropy generative model of EoR lightcones, which we quantitatively validated by comparing the new synthesised EoR lightcones with the single target lightcone from which the model is defined. Using independent statistics such as the power spectrum, histograms, and Minkowski functionals, we show that the synthesised lightcones agree very well with the target lightcone, both statistically and visually. The success of these generative models in quickly generating EoR lightcones, which can be extended to a broad range of 3D and heterogeneous 2+1D data, opens up a variety of potential applications, from forward modelling to uncertainty quantification.

**Key words.** instrumentation: interferometers – methods: statistical – cosmology: observations – early Universe – dark ages, reionization, first stars

## 1. Introduction

The Epoch of Reionisation (EoR) is a critical transition in the history of the Universe, marking the emergence of the first stars and the eventual ionisation of the once-neutral intergalactic medium (IGM). Indirect observations of the EoR from the cosmic microwave background (CMB) (Planck Collaboration XLVII 2016; Gorce et al. 2018) and the Gunn-Peterson trough in quasar spectra (Fan et al. 2006; Becker et al. 2015; Bosman et al. 2018; Qin et al. 2021), constrain the timing of the EoR between redshifts  $14 \gtrsim z \gtrsim 5.3$ .

Ground-based CMB measurements provide a complementary EoR probe: South Pole Telescope (SPT) analyses constrain the duration to  $\Delta z_{\text{re},50} = 1.1^{+1.6}_{-0.7}$  ( $<4.1$  at 95% C.L.), and a cosmology-dependent reanalysis finds  $\Delta z_{\text{re},50} < 4.3$  with  $z_{\text{re}} = 7.9^{+1.1}_{-1.3}$  (Reichardt et al. 2021; Gorce et al. 2022); although the Atacama Cosmology Telescope (ACT) has not yet detected the trispectrum, it sets an upper limit and roadmap for future detections (MacCrann et al. 2024), reinforcing the high- $\ell$  trispectrum as a viable, complementary EoR probe (Kumar et al. 2025).

While these indirect observations offer loose timings of the EoR, they do not constrain the sources responsible for the EoR, nor the state of the IGM, for which a direct probe is needed. The redshifted 21 cm signal, resulting from the hyperfine transition in neutral hydrogen during the EoR, presents a promising direct observational probe (observed by radio interferometers

between  $\sim 100$  MHz and 200 MHz) for a comprehensive study of this epoch.

Observations of the 21 cm signal (EoR signal, henceforth) by radio interferometers must contend with foregrounds such as synchrotron emission from diffuse clouds in the Galaxy, noise arising from the antennas, and systematics such as mode-mixing arising from the chromatic nature of radio interferometers. All of these contaminants are orders of magnitude stronger in total variance than the signal itself (Shaver et al. 1999; Jelić et al. 2008; Hothi et al. 2021).

The cylindrically averaged power spectrum of a radio interferometer observation (Fig. 1) has three distinct regions in the  $k_{\parallel}-k_{\perp}$  plane<sup>1</sup> (Datta et al. 2010; Vedantham et al. 2012). The first region, located at low  $k_{\parallel}$ , is dominated by astrophysical foregrounds. These foregrounds include Galactic emission, arising from free-free bremsstrahlung and diffuse synchrotron emission, as well as extragalactic synchrotron emission arising from sources such as radio haloes and clusters (Jelić et al. 2008; Hothi et al. 2021). The second region is the foreground wedge, which arises from the chromatic nature of the instrument: the interferometer's point spread function (PSF) is frequency dependent, leading to mode mixing between  $k_{\perp}$  and  $k_{\parallel}$  and causing the foreground emission to leak to higher  $k_{\parallel}$  (Liu et al. 2014). Bright sources entering the side lobes of the interferometer's PSF

<sup>1</sup> The  $z$ -domain, or line of sight, is denoted by  $\parallel$ , and the  $xy$ -plane, the perpendicular sky domain, denoted by  $\perp$ .

\* Corresponding author: Ian.Hothi@obspm.fr

also contribute to the emission within the wedge, all of which are bound by the horizon<sup>2</sup>. The third region is the EoR window, where foregrounds and instrumental effects are subdominant (there is little to no contribution on these scales) to the noise and EoR signal (see [Munshi et al. 2025](#) for a more comprehensive overview of the wedge).

The EoR signal is expected to be highly non-Gaussian, and this non-Gaussian information is important to constrain both the astrophysics of this epoch and the progression of reionisation ([Shimabukuro et al. 2017](#); [Watkinson et al. 2022](#); [Tiwari et al. 2022](#)). To fully characterise the signal requires statistics beyond the power spectrum.

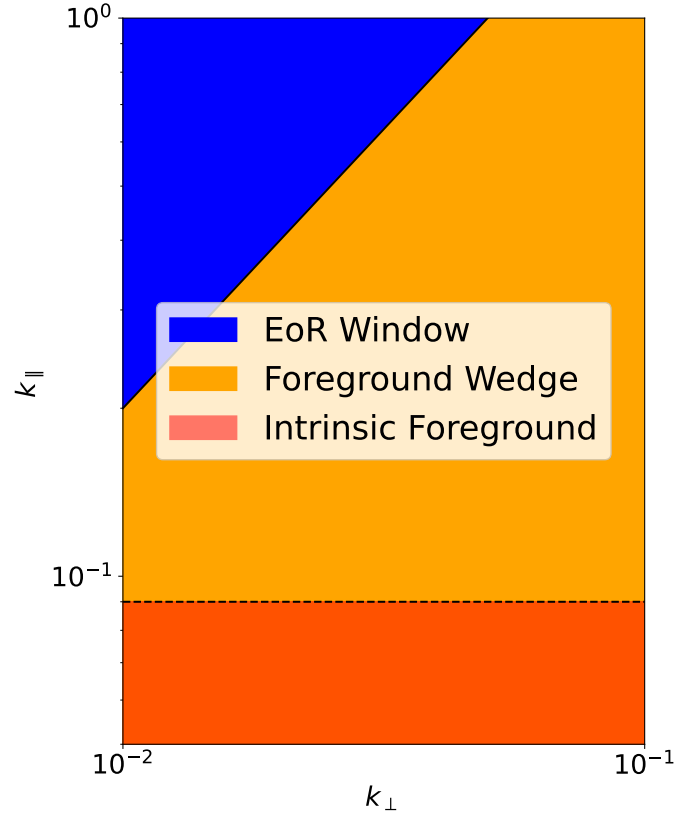
Scattering transforms (STs) have emerged as a novel set of summary statistics constructed by successive wavelet convolutions and the application of non-linear operators such as modulus ([Mallat 2011](#); [Bruna & Mallat 2012](#)). They have emerged as powerful low-variance statistics to characterise non-Gaussian processes and have already been successfully applied to various topics, including the study of the interstellar medium ([Allys et al. 2019](#); [Regaldo-Saint Blancard et al. 2020](#); [Lei & Clark 2023](#)); weak lensing ([Cheng et al. 2020](#)); the large-scale structure of the Universe for 2D ([Allys et al. 2020](#)) and 3D data ([Eickenberg et al. 2022](#); [Valogiannis & Dvorkin 2022](#); [Régaldó-Saint Blancard et al. 2024](#)); and more recently, to the EoR ([Greig et al. 2022](#); [Zhao et al. 2024](#); [Hothi et al. 2024](#); [Prelogović & Mesinger 2024](#); [Shimabukuro et al. 2025](#)).

An advantage of ST statistics is that they can be used to build generative models of physical processes within the framework of maximum entropy models. Such models, parametrised by the ST statistics themselves, efficiently approximate complex non-Gaussian physical fields, even when constructed from a single sample ([Allys et al. 2020](#); [Cheng et al. 2023](#); [Mousset et al. 2024](#); [Campeti et al. 2025](#)). Scattering transforms (STs) have also been used in the context of EoR to build diffusion models, using U-Nets, and outperform generative adversarial networks when applied to 2D EoR fields ([Zhao et al. 2023](#)).

The ability to generate realistic realisations of various processes under ST constraints has led to the development of new component separation algorithms. These ST-based component separations have already been successfully applied directly to observational data, exploiting the non-Gaussian properties of the components within the data, without physically driven priors on the component of interest ([Regaldo-Saint Blancard et al. 2021](#); [Delouis et al. 2022](#); [Auclair et al. 2024](#)). However, these previous algorithms have only been applied to 2D data. For EoR studies, this formalism must be extended to 3D EoR lightcones. This in turn requires reliance on ST statistics that can adequately describe such data, the development of which is the purpose of this paper. This work therefore extends ST generative models to 3D EoR data and quantitatively validates their quality.

To construct such 3D ST statistics, the first step is to introduce a relevant set of wavelets that characterise the different regions of the  $k_{\parallel}$ - $k_{\perp}$  plane (see Fig. 1). In previous 3D ST applications, wavelet sets typically probe extended regions of the  $k_{\parallel}$ - $k_{\perp}$  plane, thus mixing regions where the foregrounds are dominant and subdominant (see, for example, [Eickenberg et al. 2022](#)). This is problematic because mixing regions where foregrounds are dominant with subdominant regions is detrimental to the efficient extraction of information from the EoR signal ([Blamart & Liu 2025](#)). To avoid this drawback, we constructed a new set of 3D wavelets from the tensor product of 1D wavelets

<sup>2</sup> However, foregrounds can leak across the wedge line ([Murray & Trott 2018](#)).



**Fig. 1.** Schematic of the cylindrically-averaged power spectrum of an observed EoR lightcone, with the three distinct regions highlighted.

along the  $k_{\parallel}$  axis and 2D isotropic wavelets along the sky axes ( $k_x, k_y$ ). We then used this new wavelet set to construct 3D ST statistics, which we in turn used to construct a generative model of EoR lightcones, the quality of which was assessed quantitatively. The ST model is built from a single simulated EoR lightcone; we compared its statistical properties with those of the simulation using several independent and complementary statistics.

The paper is structured as follows. Section 2 outlines the simulations used. Section 3 presents the 3D scattering transforms. Section 4 describes the generative model built from scattering transforms. Section 5 compares the synthesised lightcones with the target EoR lightcones.

## 2. LoReLi II simulation

The simulation used in this work is from the LoReLi II (LOW REsolution LICorice) database ([Meriot & Semelin 2024](#); [Meriot et al. 2025](#)). This database is a collection of Licorice simulations designed to study the EoR signal ([Semelin et al. 2007, 2017](#)).

The LoReLi II database comprises 9828 simulations of the EoR signal brightness temperature and explores a four-parameter space: the gas-to-star conversion timescale, the minimum halo mass for star formation, the escape fraction of UV radiation, and the X-ray production efficiency, while keeping the cosmological parameters constant at  $\{H_0: 67.8 \text{ km/s/Mpc}, \Omega_0: 0.308, \Omega_b: 0.0484, \Omega_\Lambda: 0.692, \text{ and } \sigma_8: 0.815\}$ . This setup allows for a detailed exploration of the astrophysical processes during reionisation without the confounding effects of varying cosmology.

In this paper, the ST generative model is built from a single LoReLi simulation of a lightcone. The astrophysical parameters of the lightcone, which we refer to as the target lightcone, are the log10 of the minimum halo mass,  $M_{\min} = 8.0$ ; the X-ray luminosity fraction,  $f_x = 0.1$ ; the hard X ray fraction,  $r_{\frac{H}{S}} = 0.2$ ; the gas conversion timescale for star formation in units of 10 Myr,  $\tau = 8191.9$ ; and the ionising escape fraction,  $f_{\text{esc,post}} = 0.275$ .

This lightcone spans a redshift range of  $23 \leq z \leq 5$ . The  $xy$ -plane (the 2D sky) contains  $256^2$  pixels, and the  $z$ -domain contains 3072 redshift channels between  $z = 23$  and  $z = 5$ . To reduce the computational cost and limit the  $z$  evolution over the lightcone, we downgraded it from  $256^2$  pixels to  $128^2$  pixels in the  $xy$  plane, and restricted the  $z$  values to  $9.5 \leq z \leq 8.5$  (frequency bandwidth  $\delta\nu \sim 14$  MHz), corresponding to 170 pixels in the  $z$  domain, for a final data cube of (128, 128, 170) pixels. If the ST statistics were estimated on the initial lightcone, their statistical properties would evolve with  $z$ , whereas we approximated only limited  $z$ -evolution in the final constructed lightcone. However, we note that a lightcone with a larger redshift range could be built by stacking models constructed with a more limited redshift range (Blamart & Liu 2025).

### 3. Scattering transforms for 3D lightcones

#### 3.1. Wavelets for 3D lightcones

In this paper, we introduce a wavelet set that characterises the  $k_{\parallel}$ - $k_{\perp}$  plane by performing a tensor product between wavelet sets defined on different axes. The first is a 1D wavelet set defined in  $k_z$ , denoted by  $\parallel$ , and the second is a 2D isotropic wavelet set defined in the  $(k_x, k_y)$  plane, denoted by  $\perp$ . To simplify notation, we define the modulus wavenumbers in the two  $\parallel$  and  $\perp$  directions by  $k_{\parallel} = |k_z|$  and  $k_{\perp} = \sqrt{k_x^2 + k_y^2}$ .

The L1-normalised mother wavelet of the 1D  $\parallel$  wavelet set is defined as

$$\Psi_0^{\parallel}(k_z) = \psi_0^{\parallel}(k_{\parallel}) = \frac{1}{\sqrt{2\pi\sigma_0^{\parallel 2}}} \exp\left(-\frac{(k_{\parallel} - \mu_0^{\parallel})^2}{2\sigma_0^{\parallel 2}}\right), \quad (1)$$

where  $\mu_0^{\parallel}$  is the central wavenumber of the mother wavelet in Fourier space, and  $\sigma_0^{\parallel}$  its width. Here, we denote the wavelets in the 3D  $(k_x, k_y, k_z)$  space with  $\Psi$  and the wavelets in the 2D  $(k_{\parallel}, k_{\perp})$  cylindrical plane with  $\psi$ . The L1-normalised mother wavelet of the 2D  $\perp$  isotropic wavelet set is defined as

$$\Psi_0^{\perp}(k_x, k_y) = \psi_0^{\perp}(k_{\perp}) = \frac{1}{2\pi\sigma_0^{\perp 2}} \exp\left(-\frac{(k_{\perp} - \mu_0^{\perp})^2}{2\sigma_0^{\perp 2}}\right), \quad (2)$$

where  $\mu_0^{\perp}$  is again the central wavenumber of the mother wavelet in Fourier space, and  $\sigma_0^{\perp}$  its width.

We then constructed the two wavelet sets by dilating the respective mother wavelets,

$$\Psi_{j_{\parallel}}^{\parallel}(k_z) = \psi_{j_{\parallel}}^{\parallel}(k_{\parallel}) = 2^{-\frac{1}{2} \frac{j_{\parallel}}{Q_{\parallel}}} \psi_0^{\parallel}\left(2^{-\frac{j_{\parallel}}{Q_{\parallel}}} k_{\parallel}\right), \quad (3)$$

$$\Psi_{j_{\perp}}^{\perp}(k_x, k_y) = \psi_{j_{\perp}}^{\perp}(k_{\perp}) = 2^{-\frac{1}{2} \frac{2j_{\perp}}{Q_{\perp}}} \psi_0^{\perp}\left(2^{-\frac{j_{\perp}}{Q_{\perp}}} k_{\perp}\right). \quad (4)$$

Here,  $j_{\parallel}$  and  $j_{\perp}$  are the integer scales of the dilation, and  $Q_{\parallel}$  and  $Q_{\perp}$  are the global quality factors. From these equations,  $\Psi_{j_{\parallel}}^{\parallel}(k_z)$  and  $\Psi_{j_{\perp}}^{\perp}(k_{\perp})$  can be written in the same form as the mother

wavelets defined in Eqs. (1) and (2) by making the following replacements:

$$\begin{aligned} \sigma_0^{\alpha} &\rightarrow \sigma_{j_{\alpha}}^{\alpha} = 2^{\frac{j_{\alpha}}{Q_{\alpha}}} \sigma_0^{\alpha}, \\ \mu_0^{\alpha} &\rightarrow \mu_{j_{\alpha}}^{\alpha} = 2^{\frac{j_{\alpha}}{Q_{\alpha}}} \mu_0^{\alpha}, \end{aligned} \quad (5)$$

for  $\alpha \in \{\parallel, \perp\}$ . The  $j_{\alpha}$  integer scales of the  $\psi_{j_{\alpha}}$  wavelets range from 0 to  $J_{\alpha} - 1$ , where  $J_{\alpha}$  is the number of scales. Once  $J_{\alpha}$  is chosen, the value of  $Q_{\alpha}$  is fixed so that the wavelet set samples the full extent of the Fourier space. Specifically, when the mother wavelet, centred at  $\mu_0^{\alpha}$ , is dilated up to the largest scale, the resulting  $\psi_{J_{\alpha}-1}$  wavelet is centred at the maximum Fourier scale  $k_{\alpha,\text{max}}$ , i.e. the Nyquist scale, yielding

$$\log \frac{k_{\alpha,\text{max}}}{\mu_0^{\alpha}} = \frac{J_{\alpha} - 1}{Q_{\alpha}} \log 2. \quad (6)$$

To construct the 3D wavelet, we consider the tensor product between the two wavelet sets we previously defined:

$$\begin{aligned} \Psi_{j_{\parallel}, j_{\perp}}^{\parallel}(k_x, k_y, k_z) &= \Psi_{j_{\parallel}}^{\parallel}(k_z) \otimes \Psi_{j_{\perp}}^{\perp}(k_x, k_y) \\ &= \psi_{j_{\parallel}, j_{\perp}}^{\parallel}(k_{\parallel}, k_{\perp}) = \psi_{j_{\parallel}}^{\parallel}(k_{\parallel}) \cdot \psi_{j_{\perp}}^{\perp}(k_{\perp}), \end{aligned} \quad (7)$$

where  $\otimes$  is the tensor product. In the following, we use the shorthand notation  $\lambda = \{j_{\parallel}, j_{\perp}\}$ , which yields

$$\Psi_{\lambda}(k_x, k_y, k_z) = \psi_{\lambda}(k_{\parallel}, k_{\perp}). \quad (8)$$

These wavelets are real in Fourier space and complex in real space.

In this paper, the central frequencies  $\{\mu_0^{\parallel}, \mu_0^{\perp}\}$  of the mother wavelets were chosen to be twice the smallest Fourier scale available, giving  $\{\mu_0^{\parallel}, \mu_0^{\perp}\} = \{0.023 h \text{ Mpc}^{-1}, 0.028 h \text{ Mpc}^{-1}\}$ . This choice allows sampling on the largest possible scales, while still enabling the estimation of statistics on non-periodic fields (see Sect. 4.4). We then chose  $(J_{\parallel}, J_{\perp}) = (6, 6)$ , corresponding to  $(Q_{\parallel}, Q_{\perp}) = (0.90, 0.93)$ , which is close to the commonly used dyadic scaling, resulting in a total of 36 wavelets. Finally, we chose the values  $\{\sigma_0^{\parallel}, \sigma_0^{\perp}\} = \{0.011 h \text{ Mpc}^{-1}, 0.015 h \text{ Mpc}^{-1}\}$  to cover the entire  $k_{\parallel}$ - $k_{\perp}$  Fourier plane as uniformly as possible. Fig. 2 shows the resulting wavelet set in this cylindrical Fourier space.


This wavelet set applies not only to EoR lightcones but also to all spectroscopic datasets and to any 2+1D field whose third axis represents a different quantity. In this work, we demonstrate its performance on EoR lightcones with limited redshift evolution ( $\Delta z \simeq 1$ ); extending the method to fully evolving lightcones with larger  $\Delta z$  is a direction for future work.

#### 3.2. Scattering transform statistics

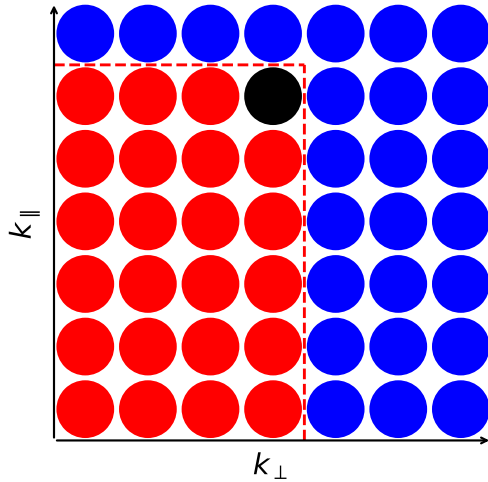
Scattering transform (ST) statistics are a family of summary statistics. In this work, we consider a simplified version<sup>3</sup> of the scattering covariance statistics introduced in Cheng et al. (2023). Consider a 3D lightcone  $I$ . Its scattering covariance statistics are defined as

$$C_{\lambda_1, \lambda_2}^{p_1, p_2} = \text{Cov}([I \otimes \Psi_{\lambda_1}]^{p_1}, [I \otimes \Psi_{\lambda_2}]^{p_2}), \quad (9)$$

<sup>3</sup> Compared with the statistics defined in Cheng et al. (2023), no second convolution is explicitly performed. These statistics can also be viewed as a compact form of the wavelet phase harmonics statistics introduced in Allys et al. (2020), where only  $p = 0$  and  $p = 1$  phase harmonics are considered, and no translations are performed prior to covariance estimation.



**Fig. 2.** Different regions in the  $k_{\parallel}$  vs  $k_{\perp}$  plane characterised by wavelets in the 3D wavelet set.



**Fig. 3.** Schematic of the secondary convolution criterion  $\lambda_1 \leq \lambda_2$ , for the 3D wavelets set in the  $k_{\parallel}$  vs  $k_{\perp}$  plane. The starting wavelet  $\psi_{\lambda_1}$  (black) and the second wavelet  $\psi_{\lambda_2}$  probe regions such that  $\mu_{j_{\parallel,2}}^{\parallel} \leq \mu_{j_{\parallel,1}}^{\parallel}$  and  $\mu_{j_{\perp,2}}^{\perp} \leq \mu_{j_{\perp,1}}^{\perp}$ . These wavelets, for the example  $\psi_{\lambda_1}$ , are shown in the highlighted region (red).

where  $\otimes$  denotes a convolution and

$$[z]^p = \begin{cases} |z|, & \text{if } p = 0. \\ z, & \text{if } p = 1. \end{cases} \quad (10)$$

The covariance between two stochastic processes  $X$  and  $Y$  is defined as  $\text{Cov}(X, Y) = \langle XY^* \rangle - \langle X \rangle \langle Y^* \rangle$ , where  $\langle \rangle$  denotes the expected value and  $*$  denotes the complex conjugate. In this work, the expected values were estimated using a spatial average.

Using this definition, the three subsets of scattering covariance coefficients are:

$$C^{11} = \text{Cov}(I \otimes \Psi_{\lambda_1}, I \otimes \Psi_{\lambda_2}), \quad \lambda_1 = \lambda_2, \quad (11)$$

$$C^{01} = \text{Cov}(|I \otimes \Psi_{\lambda_1}|, I \otimes \Psi_{\lambda_2}), \quad \lambda_1 \leq \lambda_2, \quad (12)$$

$$C^{00} = \text{Cov}(|I \otimes \Psi_{\lambda_1}|, |I \otimes \Psi_{\lambda_2}|), \quad \lambda_1 \leq \lambda_2. \quad (13)$$

In these equations, the condition  $\lambda_1 \leq \lambda_2$  denotes  $\mu_{j_{\parallel,2}}^{\parallel} \leq \mu_{j_{\parallel,1}}^{\parallel}$  and  $\mu_{j_{\perp,2}}^{\perp} \leq \mu_{j_{\perp,1}}^{\perp}$ , as summarised in Fig. 3. This  $\lambda$ -scale conditioning is explained as follows. For  $C^{11}$ , the covariance is zero if  $\lambda_1 \neq \lambda_2$ . Since the spectral supports of different wavelets do not overlap,  $I \otimes \Psi_{\lambda_1}$  and  $I \otimes \Psi_{\lambda_2}$  also have distinct spectral supports and therefore exhibit vanishing covariance, as they are linearly decorrelated<sup>4</sup>. For  $C^{01}$ , convolving  $I \otimes \Psi_{\lambda_1}$  with the first wavelet concentrates the resulting field into a band of scales around  $\lambda_1$ . Applying the modulus operator,  $|I \otimes \Psi_{\lambda_1}|$ , demodulates the field by re-centering its support towards lower scales (see, e.g. Zhang & Mallat 2019). To obtain a non-vanishing covariance, the secondary wavelet convolution  $I \otimes \Psi_{\lambda_2}$ , centred at scales  $\lambda_2$ , must now overlap the envelope defined by  $|I \otimes \Psi_{\lambda_1}|$ , enforcing  $\lambda_1 \leq \lambda_2$ . For  $C^{00}$ , since the covariance is commutative,  $|I \otimes \Psi_{\lambda_1}|$  and  $|I \otimes \Psi_{\lambda_2}|$  play the same role, so the condition  $\lambda_1 \leq \lambda_2$  avoids double counting.

Although the scattering covariances  $C^{p_1, p_2}$  capture the non-Gaussian properties of a given lightcone, they also depend on the amplitude of its power spectrum. To correct for this dependence, we normalised the scattering covariances by the  $C^{11}$  of a reference field. In this work, this reference field is the target lightcone and the normalised scattering covariance statistic is denoted  $\bar{C}_{\text{ref}}^{11}$ . The normalised scattering covariance statistics are defined as follows:

$$\bar{C}^{00}(\lambda_1, \lambda_2) = \frac{C^{00}(\lambda_1, \lambda_2)}{\sqrt{C_{\text{ref}}^{11}(\lambda_1) C_{\text{ref}}^{11}(\lambda_2)}}, \quad (14)$$

$$\bar{C}^{01}(\lambda_1, \lambda_2) = \frac{C^{01}(\lambda_1, \lambda_2)}{\sqrt{C_{\text{ref}}^{11}(\lambda_1) C_{\text{ref}}^{11}(\lambda_2)}}, \quad (15)$$

$$\bar{C}^{11}(\lambda) = \frac{C^{11}(\lambda)}{C_{\text{ref}}^{11}(\lambda)}. \quad (16)$$

To further constrain the power spectrum, we used an additional set of  $C^{11}$  coefficients, denoted  $C^{11'}$ . These coefficients were computed according to Eq. (11), with a set of wavelets defined as in Sect. 3.1, but with a different number of scales,  $J'_\alpha > J_\alpha$ . These  $C^{11'}$  coefficients, computed from a larger set of wavelets, used in conjunction with  $C^{11}$ , provide an overall improved spectral resolution for power spectrum constraints (Mousset et al. 2024), without increasing the computational complexity for higher-order terms. For these coefficients, we used  $(J'_\parallel, J'_\perp) = (15, 15)$ , corresponding to  $(Q_\parallel, Q_\perp) = (2.53, 2.60)$ .

We concatenated the various ST coefficients to form the set of ST statistics:

$$\phi(I) = \{\bar{C}^{11}, \bar{C}^{11'}, \bar{C}^{00}, \bar{C}^{01}\}. \quad (17)$$

The choice of parameters to define the wavelet sets results in 36 coefficients in  $C^{11}$ , 225 coefficients in  $C^{11'}$ , and 441 coefficients in  $C^{01}$  and  $C^{00}$ , for a total of 1143 ST statistics.

<sup>4</sup> This can be shown using Plancherel's Theorem.

## 4. Generative models

### 4.1. Maximum entropy microcanonical models

We constructed a maximum entropy model of an EoR lightcone from the scattering covariance statistics on the target lightcone,  $d_t$ . The new approximate realisations of this lightcone were then sampled using gradient descent (Bruna & Mallat 2018; Allys et al. 2020; Régaldou-Saint Blancard et al. 2023). This gradient descent sampling transforms a white Gaussian distribution into the maximum entropy distribution conditioned on the ST statistics of  $d_t$ .

The procedure is as follows. The map on which the gradient descent is performed is denoted  $u$ . It is initialised with a white noise realisation,  $u = u_0$ , with a mean and standard deviation matching those of the target lightcone. The optimisation is performed by constraining the scattering transform statistic,  $\phi(u)$ , of  $u$  with the following loss:

$$\mathcal{L}(u) = \|\phi(d_t) - \phi(u)\|^2, \quad (18)$$

where  $\|\cdot\|$  is the Euclidean norm and  $d_t$  is fixed. This gradient descent is performed until the loss converges. The samples,  $\tilde{d}$ , from the ST generative model correspond to the map  $\tilde{d} = u_f$  obtained at the end of the optimisation, whose ST statistics are close to those of  $d_t$ . Different  $\tilde{d}$  samples are obtained by initialising the optimisation with different white noise realisations.

### 4.2. Numerical experiment and statistical validation

To validate the quality of the ST generative models, we compared the statistical properties of the generated lightcones, after performing the inverse quantile transform, with those of the target lightcone,  $d_t$ . We generated 30 independent synthesised lightcones, denoted  $\{\tilde{d}_i\}_{i=1\dots 30}$  after the inverse quantile transform, each initialised from a different white Gaussian noise realisation. To perform the optimisation, we computed the loss function gradient using automatic differentiation with the L-BFGS-B algorithm (Byrd et al. 1995) (as implemented in PyTorch). We performed the optimisation using 2000 gradient descent iterations, at which point the loss had converged to a plateau after being reduced by approximately ten orders of magnitude. Each optimisation for a single lightcone required two hours on a Nvidia Tesla A100 GPU with 80 GB of memory.

Beyond an initial visual comparison, the statistics considered were the histogram, power spectrum, and Minkowski functionals. In each case, we compared the mean statistics of  $d_t$  (shown in blue in comparison figures) and of the 30  $\tilde{d}_i$  (shown in green). To assess potential model bias, we compared the difference between these mean statistics and their sample standard deviation,  $\sigma_t^{\text{oct}}$ , between octants of the lightcone<sup>5</sup>, estimated on the target lightcone,  $d_t$ . The standard deviation of the St statistics estimator used to construct the generative model is  $\sigma_t^{\text{oct}}$ , since only a volume corresponding to a single octant of the lightcone was used, due to the issue of non-periodic boundary conditions, as explained above. This sampling variance can then simply explain modelling biases of this order.

Additionally, we compared  $\sigma_t$  and  $\sigma_s$ , the sample standard deviation across lightcones, estimated from  $d_t$  and the 30  $\tilde{d}_i$ , respectively. This allowed us to quantify whether the ST generative model reproduces both the mean and sample variance of the

<sup>5</sup> To create octants, the data cube was divided into eight equal subvolumes, of dimension (64, 64, 85) pixels. These octants have approximately the same volume as the region used to estimate the ST statistics due to non-periodic boundary conditions.

**Table 1.** Overview of the defined  $\sigma$  quantities.

Symbol	Definition
$\sigma_{\text{oct},t}$	Sample standard deviation across the eight octants of the target lightcone $d_t$ .
$\sigma_t$	Approximate lightcone-to-lightcone standard deviation for $\mathbf{d}_t$ , given by $\sigma_t \approx \sigma_{\text{oct},t} / \sqrt{8}$ .
$\sigma_s$	Lightcone-to-lightcone standard deviation from 30 independent synthesised lightcones $\{\tilde{\mathbf{d}}_i\}$ .
$\sigma_g$	Lightcone-to-lightcone standard deviation from 30 Gaussian realisations, having the same power spectrum as $d_t$ .

statistics. While  $\sigma_s$  can be estimated directly from the 30  $\tilde{d}_i$ ,  $\sigma_t$  has to be estimated from  $d_t$  only. To do this, we approximated the octants of  $d_t$  as independent, resulting in  $\sigma_t \approx \sigma_t^{\text{oct}} / \sqrt{8}$ . While this is a rough approximation, which will a priori tend to underestimate  $\sigma_t$ , since the octant are not independent, it enables a first comparison between  $\sigma_t$  and  $\sigma_s$ .

To provide additional comparisons, we also considered the distributions of statistics obtained from 30 Gaussian realisations with the same power spectrum as  $d_t$  (shown in orange), with their associated lightcone-to-lightcone sample standard deviation,  $\sigma_g$ , and with the distributions obtained from the complete set of lightcones in the Loreli II database (shown in purple).

A list of the different  $\sigma$  values defined in this section is given in Table 1.

### 4.3. Data conditioning

Prior to the construction of the generative model, the data was conditioned to mitigate several practical issues that could adversely affect the resulting synthesis. This data conditioning ensured that the generative model operated on a well-suited representation of the EoR lightcone, thus improving the fidelity and stability of the subsequent synthesis.

#### 4.3.1. Quantile function

The EoR lightcone is a complex non-Gaussian process, with a highly skewed histogram that is difficult to recover using an ST maximum entropy model. To improve the modelling, we first applied an invertible non-linear transform to the EoR lightcone to obtain a histogram closer to a Gaussian distribution. We then sampled the ST maximum entropy model in this transformed data space, before returning to the original EoR lightcone data space using the inverse non-linear transform.

Such ST modelling based on an invertible non-linear transform has been used in previous studies, for example using a logarithmic pointwise transform before constructing ST generative models (Allys et al. 2020; Régaldou-Saint Blancard et al. 2023; Mousset et al. 2024). This approach is particularly useful for modelling processes with complex histograms, such as those with only non-negative values, including the matter density of large-scale structures in the Universe. The general idea of transforming a process into another one that is better approximated is also used when constructing lognormal models for processes whose logarithm is expected to be much better described by a Gaussian model than the process itself. Lognormal models are

used, for example, to describe the thermal emission of Galactic molecular clouds (Miville-Deschênes et al. 2007; Levrier et al. 2018).

In this work, we used a quantile transformation to map the histogram of the target lightcone to a normal distribution (Mark Beasley & Erickson 2009; McCullagh & Tressoldi 2020). To do this, we first estimated the global empirical cumulative distribution function (CDF)  $\hat{F}$  of the field for the entire ( $\delta\nu \sim 14$  MHz) lightcone. We then computed the corresponding percentile  $p$  for each data point  $d(\mathbf{r})$  in the lightcone,

$$p(\mathbf{r}) = \hat{F}(d(\mathbf{r})). \quad (19)$$

Next, this percentile  $p$  was passed into the inverse of the standard normal CDF (also known as the quantile function) to obtain

$$\bar{d}(\mathbf{r}) = \Phi^{-1}(p(\mathbf{r})). \quad (20)$$

This procedure produces a set of  $\bar{d}$  values that follow a normal distribution, replacing the original skewed distribution while preserving the ordering of the original data points.

In this work, the quantiles for the transform were estimated from all data points within the target lightcone. The choice of the number of quantiles corresponds to a standard bias-variance trade-off: too few quantiles result in a poor representation of the data distribution, while too many can introduce overfitting and unnecessary computational overhead. We empirically find that 10 000 quantiles provide a good balance. Given that the target lightcone contains approximately  $10^6$  pixels, using 10 000 quantiles corresponds to approximately 300 pixels per quantile. This quantile-to-pixel ratio is sufficient to accurately represent the low density tail of the distribution without incurring needless computational cost.

It should be noted that the quantile transform is not a point-wise operator but is instead applied by analysing the entire lightcone simultaneously. Although this transform efficiently models complex histograms, any modelling error in the transformed data space can cause significant instabilities in the inverse transform. In particular, the significant effect of large scales on the histogram, which are difficult to model statistically from a small amount of data, can make it difficult to use such a transform.

#### 4.3.2. Treatment of unconstrained scales

The generative model in this work does not constrain two limiting ranges of scales. The first unconstrained range, denoted  $\kappa_{\text{low}}$ , corresponds to the largest scales, which have Fourier mode wavenumbers smaller than the mother wavelet's wavenumber,  $\mu_0$ . The second unconstrained range, denoted  $\kappa_{\text{high}}$ , corresponds to the smallest scales, which have Fourier modes beyond the Nyquist frequency.

Since we first constructed the ST generative model on the quantile-transformed data before applying the inverse quantile transform, these unconstrained scales can significantly affect the final synthesised lightcone because the transform is the non-linear and non-local. To mitigate their impact, we handled them explicitly, as described below:

- The  $\kappa_{\text{low}}$  scales were replaced with the corresponding modes from the target lightcone in the transformed data space. This substitution occurred prior to the gradient descent optimisation, ensuring that these large-scale modes were included in the initial conditions. This was necessary because, if left unconstrained, these scales would influence the global inverse quantile transform. Since these scales are not characterised by the ST statistics, they do not directly influence the

generation of the remaining scales during the optimisation process.

- The  $\kappa_{\text{high}}$  scales in the target lightcone were filtered prior to the quantile transform, both when constructing the generative model and when comparing with the synthesised maps. Although the non-linearity of the quantile transform introduces some beyond-Nyquist scales after the histogram transform, their effect on the inverse histogram transform after modelling is negligible.


In principle, our framework can be run without re-inserting the modes with  $k < k_{\text{low}}$  from the target lightcone. In this case, these modes remain unconstrained by the optimisation process. The generative model still visually reproduces the target lightcone and accurately matches the scattering statistics on intermediate and small scales. However, the largest scales show biased power, with deviations at  $k \lesssim \kappa_{\text{low}}$  comparable to, although generally within, the sample variance of the target lightcone. We also experimented with replacing the modes with Gaussian realisations matched to the target power spectrum. However, because the quantile transform is a global, non-linear mapping, the large-scale modes strongly influence the histogram. Consequently, even modest higher-order mismatches propagate into systematic biases at low  $k$  after the inverse transform. A more physical alternative would be to source these large-scale modes from a dedicated low-resolution EoR simulation. However, consistently coupling them to the high-resolution, quantile-transformed small scales is non-trivial, and we leave this as an interesting direction for future work.

#### 4.4. Non-periodic boundary conditions


Simulations of the EoR lightcone have non-periodic boundaries, in particular because of the redshift evolution along the line of sight. However, since we compute convolutions by products in Fourier space, these are inherently periodic. To avoid taking boundary discontinuities into account when estimating the ST statistics, the empirical covariances were computed only over a central smaller cubic volume of the data<sup>6</sup>. This is achieved by cropping the edges of the wavelet convolved fields, so that all pixels over which the average is computed are sufficiently separated from the boundaries. To verify this, the cropped band sizes are matched to the maximum sizes characterised by the wavelets, which are determined by the mother wavelet scale  $\frac{1}{\mu_0}$ . Specifically, we removed a width at each end of each axis equal to half the mother wavelet scale on that axis, yielding an average over a volume of (61, 61, 87) pixels.

Note that while the statistics of the target lightcones, estimated on  $d_i$  and used to constrain the generative model, were computed taking into account the non-periodic boundary conditions as described above, the samples generated by this generative model are full lightcones with periodic boundary conditions. To achieve this, their initial conditions were white noise defined over the entire (128, 128, 170) pixels data cube, and their statistics were computed using period convolutions and averages over the entire cube. We performed these syntheses by comparing the two types of ST statistics estimators, with and without periodic boundary conditions, in Eq. (18). As a side effect, the statistics used to construct the generative models were estimated from only about one eighth of the target lightcone volume, leading to a higher sampling variance.

<sup>6</sup> In practice a rectangular parallelepiped, as the original lightcone.



**Fig. 4.** Comparison of a given  $xy$ -slice (single frequency channel in the lightcone), for the target EoR lightcone and three synthesised lightcones generated using the 3D scattering transforms with the inverse quantile transform applied. The synthesised lightcones visually reproduce the target lightcone well.



**Fig. 5.** Same as Fig. 4, but for a given  $xz$ -slice (with  $z$  representing the redshift domain) for the target EoR lightcone and three synthesised lightcones. As in Fig. 4, the target lightcone is visually well reproduced by the synthesised lightcones.

## 5. Validation of the generative models

### 5.1. Visual validation

A first test of synthesis quality is to visually compare three realisations,  $\tilde{d}_i$ , of a synthesised lightcone with the target lightcone,  $d_t$ . The  $xy$ -slice comparison (Fig. 4) shows that these three realisations are visually similar to  $d_t$ . The target lightcone contains several distinctive features, such as strongly localised absorbing (negative) regions and ionised (near-zero) regions, which serve as touchstones that the synthesised lightcones successfully capture. Fig. 5 shows that the synthesised lightcones again reproduce the different phases of the field (emission, absorption, and ionised) and the elongated structures characteristic of the redshift axis. However, the generated samples appear to exhibit a greater diversity in the distribution of ionised regions, although a quantitative assessment is difficult from a single  $d_t$  lightcone.

### 5.2. Histogram

Fig. 6 compares normalised histograms (i.e. the estimated probability density functions) of the target (blue) and synthesised (green) lightcones. The target lightcone exhibits a non-symmetric, left-skewed tail distribution, a clear, highly non-Gaussian signature, which the synthesised lightcones accurately reproduce over approximately four orders of magnitude, down to  $\delta T_b = -7$  mK. The bottom plot of Fig. 6 shows that the difference between these histograms, normalised by  $\sigma_t^{\text{oct}}$ , mostly lies between  $-1$  and  $1$ , consistent with the estimation variance of the target ST statistics. The top plot of Fig. 6 shows that  $\sigma_s$ , the lightcone-to-lightcone standard deviation for the synthesised maps, is of the same order as the bin-to-bin variability of the histogram estimated on  $d_t$ , indicating that this variance is also well

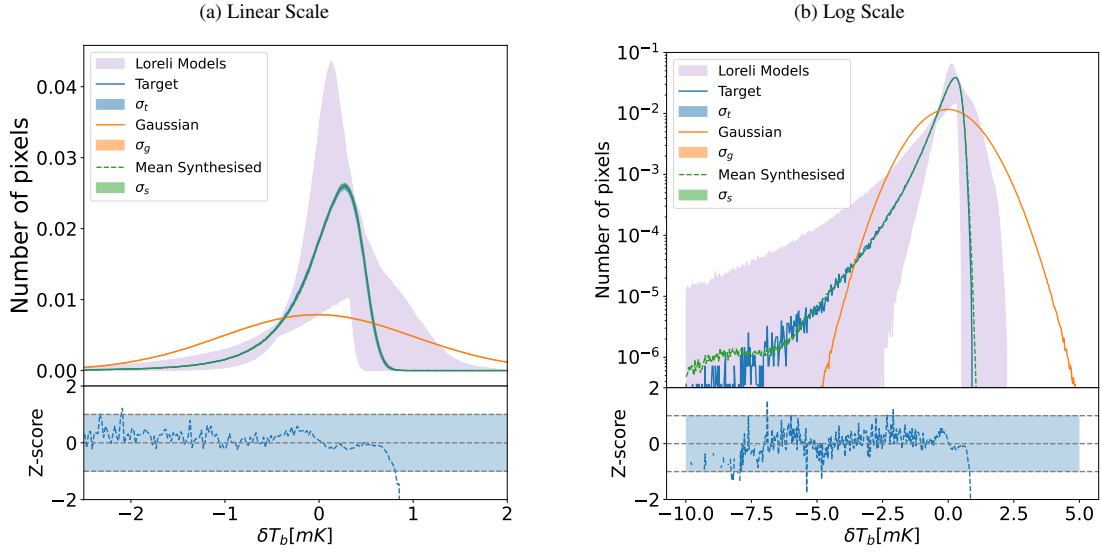
reproduced. The standard deviation,  $\sigma_t$ , is not shown in the top plot, as the sample variability is already apparent due to the fine binning.

We used negative brightness temperature values ( $\delta T_b$ ) as a proxy for ionised pixels to infer the neutral fraction of the lightcone. The neutral fraction of the target lightcone is  $0.617$ . The mean neutral fraction of the 30 generated lightcones is  $0.624 \pm 0.001$ , indicating that the generative model can recover the target lightcone's neutral fraction to 1% accuracy.

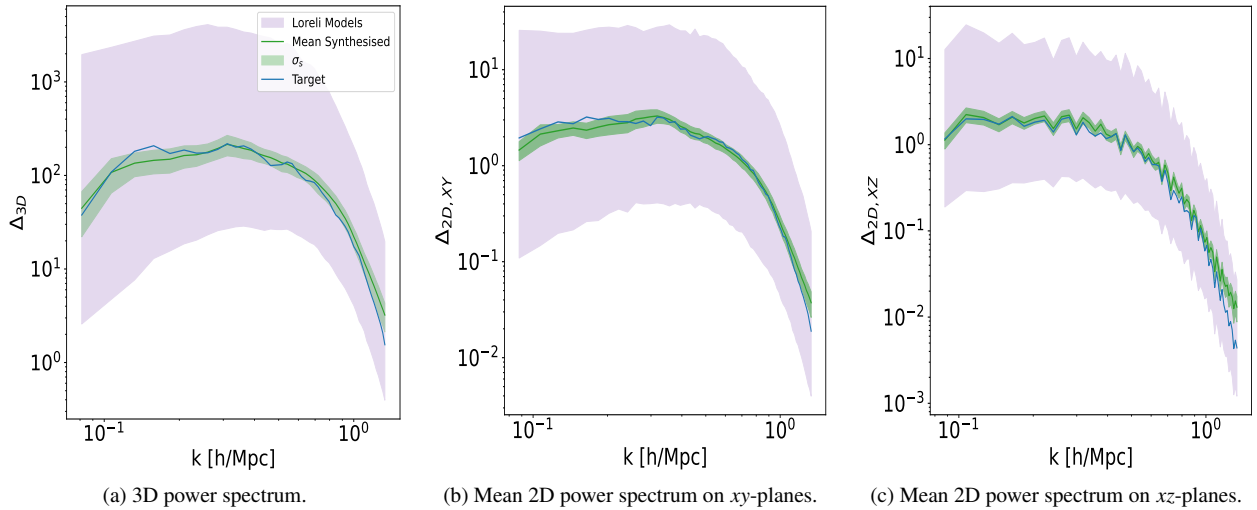
### 5.3. Power spectrum

Four power spectra were used to compare the target and synthesised lightcones: the 3D isotropic power spectrum; the 2D isotropic power spectrum on the  $xy$  or  $xz$  planes; and the cylindrically averaged power spectrum. Although the cylindrically averaged power spectrum is sometimes referred to as the 2D power spectrum, here it specifically refers to the power spectrum estimated on actual 2D fields: each 2D  $xy$  or  $xz$  slice, averaged over the  $z$  or  $y$  axis, respectively. We refer the reader to Appendix A for more details. Comparisons with Gaussian realisations are not included in this section, as these have the same power spectrum as the target lightcone.

Figure 7a compares the 3D isotropic power spectrum of the target lightcone and that of the synthesised lightcones. The power spectrum of the synthesised lightcones closely follows that of the target lightcone, remaining within  $\pm\sigma_s$  on most scales. Although estimating  $\sigma_t^{\text{oct}}$  for the power spectrum is challenging, due to the differing numbers of modes per bin for different data sizes, the averaged 3D power spectrum of the synthesised lightcones reproduces that of  $d_t$  within 25% across the range  $0.1$  and  $1 h \text{ Mpc}^{-1}$ . Figures 7b and 7c show sim-



**Fig. 6.** Linear- and log-scale histograms of the target EoR lightcone and a synthesised lightcone after the inverse quantile transform. The histogram of the synthesised lightcone, showing the mean of 30 realisations and their standard deviation (shaded region), recovers the histogram of the target lightcone well, including its skewness.




**Fig. 7.** Comparison between the target and synthesised lightcones after the inverse quantile transform, showing power spectra across different dimensions. (a) Spherically-averaged 3D power spectrum. (b) 2D power spectrum estimated on  $xy$ -planes. (c) 2D power spectrum estimated on  $xz$ -planes. In all cases the power spectra of the synthesised lightcones reproduce the largest scales well, with increasing deviations at smaller scales.

ilar results for the 2D power spectra in the  $xy$ - and  $xz$ -planes. In these cases, a visible bias appears at the smallest scales, with syntheses reproducing the  $d_t$  power spectra to within 15% on average across  $0.2-1 h \text{Mpc}^{-1}$ . The shape of the power spectra in Figure 7c results from discretisation effects introduced by the binning scheme in the cylindrical Fourier plane, which imposes a structure on the spectra.

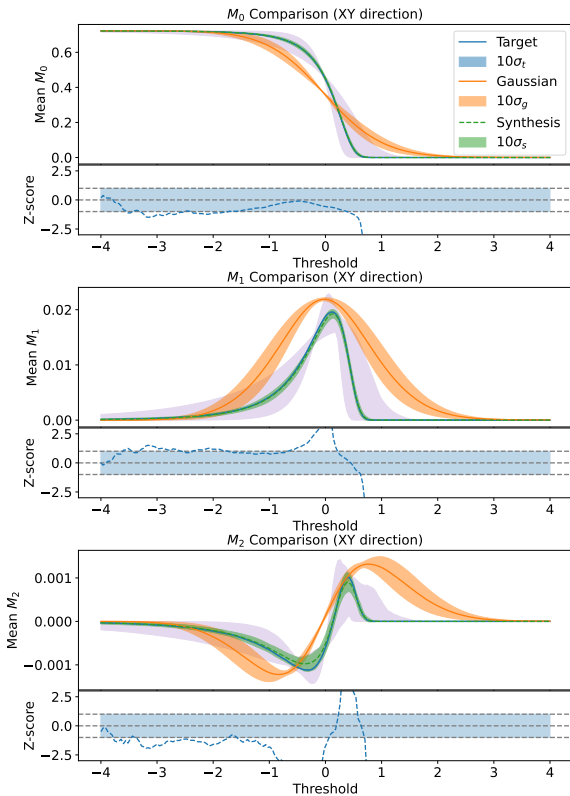
Fig. 8 compares the cylindrically averaged power spectra, where the synthesised lightcones are represented by the mean spectrum of the 30 realisations. The right panel shows the ratio of the target lightcone's cylindrically averaged power spectrum to that of the mean of the synthesised lightcones. Overall, the cylindrical power spectra are well reproduced, but Fig. 8c highlights large ratio values at high  $k_{\parallel}$  and  $k_{\perp}$ . These correspond to the bias we see at high  $k$  in the power spectrum comparison in Fig. 7, close to Nyquist frequencies. Apart from the bias observed at

high  $k$  modes, the cylindrical power spectra of the synthesised lightcones remain within  $\pm 1.5\sigma$  of the variance across the 30 realisations at all other scales.

The high- $k$  bias in the power spectra of the synthesised lightcones, particularly along the  $z$  domain, is not a limitation of the  $C^{11}$  and  $C^{11'}$  power spectrum constraints. Prior to the application of the inverse quantile transform, power spectra recover well at these scales, with only small bias. However, this bias increases significantly when the transform is subsequently applied. This likely arises because non-Gaussian structures, i.e. the statistical dependence between the different scales, are not properly reproduced at small scales. In this scenario, the highly non-linear inverse histogram transform could introduce such a bias. To address this issue, the  $C^{11}$  and  $C^{11'}$  power spectrum constraints could be added after this transform, at the expense of additional computational time.



**Fig. 8.** Comparisons of the cylindrically averaged power spectrum of the target lightcone (*left*) and synthesised lightcones (*middle*) after the inverse quantile transform. The ratio between the two cylindrically averaged power spectra is also shown (*right*), with the power spectra of the synthesised lightcones reproducing that of the target lightcone within 25%, on most scales below  $k_{\parallel} = 0.3 h \text{Mpc}^{-1}$ .



**Fig. 9.** Comparison of the Minkowski functionals applied to each frequency channel of the lightcone between the synthesised lightcones (after the inverse quantile transform) and the target lightcone. The mean Minkowski functionals across the lightcone are shown, with shaded regions indicating their standard deviation. From the bottom panel,  $M_0$  and  $M_1$  of the synthesised lightcones recover the target lightcone’s Minkowski functionals well, lying within the octant variance,  $\sigma_i^{\text{oct}}$ , and reproducing its morphology. The  $M_2$  of the synthesised lightcone follows that of the target lightcone and lies within  $\pm 2.5\sigma_i^{\text{oct}}$ .

#### 5.4. Minkowski functionals

Finally, we compared the Minkowski functionals (MFs) of the target and the synthesised lightcones, which were computed using the Minkowski functionals function of the QuantImPy Python package (Boelens & Tchelepi 2021). As with the 2D isotropic power spectrum, we computed the 2D MFs on the  $xy$  and  $xz$  planes, which were then averaged over the  $z$  or  $y$  axes,

respectively. The three MFs are  $M_0$ ,  $M_1$ , and  $M_2$  (Mecke et al. 1994; Schmalzing & Buchert 1997). These 2D MFs were computed by progressively increasing a threshold value and evaluating the properties of the regions above this threshold: total area ( $M_0$ ), perimeter ( $M_1$ ), and genus ( $M_2$ ), which provides insight into the connectivity and complexity of ionised structures.

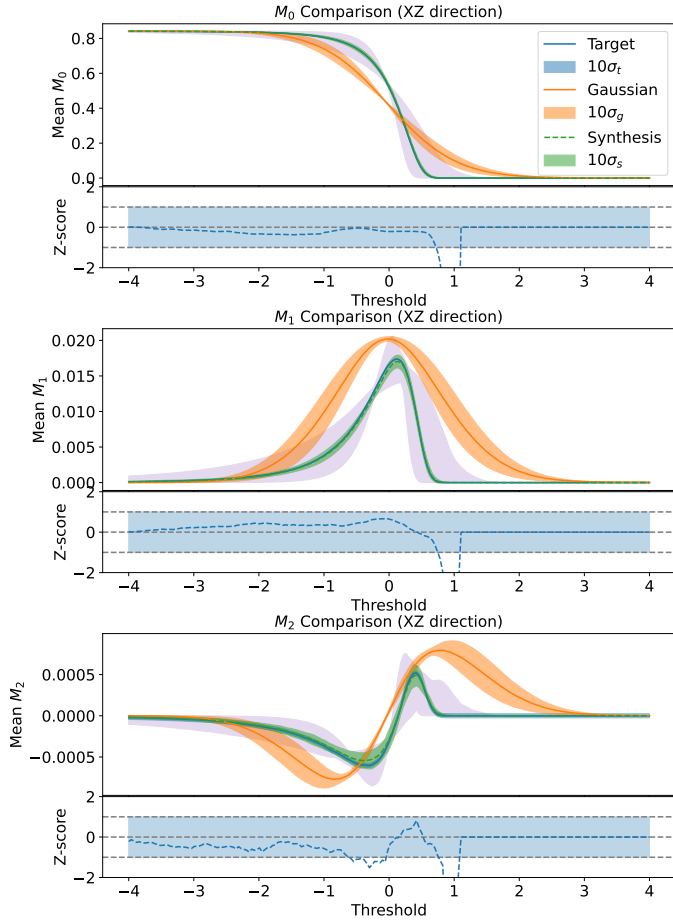
To address the non-periodic boundary conditions of  $d_t$ , the MFs were computed on the  $xy$  or  $xz$  planes after applying a complete symmetric padding to restore periodicity without significantly altering the underlying data (this results in some sharp edges at the initial boundary position). We also applied this symmetric padding when estimating the MFs of the syntheses, even though they have periodic boundary conditions, to ensure the best possible comparison.

Figures 9 and 10, show the MFs in the  $xy$  and  $xz$  planes, respectively. Since the lightcone-to-lightcone standard deviations ( $\sigma_t$ ,  $\sigma_s$ , and  $\sigma_g$ ) are very small, they are multiplied by a factor of 10 in the plots. These plots show that all three MFs of the synthesised lightcones closely resemble those of  $d_t$  in both  $xy$  and  $xz$  planes. This indicates that the ST generative models effectively reproduce the morphological features of the data. However, while the differences between these mean statistics mostly remain within  $\pm\sigma_i^{\text{oct}}$  for the  $xz$ -MFs, this is not entirely the case for the  $xy$ -MFs, for which clear deviations are particularly visible for  $M_2$ . We interpret this as the first limitation of the ST generative model, arising from its difficulty in reproducing the sharp, high-contrast, and highly textured circular features present in the target lightcone (see Fig. 4). Overall, for both  $xy$  and  $xz$  cases, the  $M_0$ ,  $M_1$ , and  $M_2$  statistics are reproduced within 1%, 10%, and 15%, respectively.

Notably, the synthesised data have a lightcone-to-lightcone standard deviation,  $\sigma_s$ , which is a few times smaller than the standard deviation  $\sigma_t$  estimated from  $d_t$ . In contrast, Gaussian models exhibit considerably larger variance. This could indicate a second limitation of ST generative models. However, care should be taken when interpreting this result, as estimating  $\sigma_t$  from  $\sigma_i^{\text{oct}}$  could lead to a substantial underestimation of this variance, which could be larger between lightcones than within them.

## 6. Conclusions

The 21 cm signal from the Epoch of Reionisation (EoR) is observed by radio interferometers as a lightcone that includes a 2D sky plane and a 1D redshift domain that captures the signal’s



**Fig. 10.** Same as Fig. 9, but for the  $xz$ -plane. The synthesised lightcones reproduce the functionals of the  $xz$ -plane, recovering the morphology of the evolving lightcone. These functionals lie within the octant variance,  $\sigma_i^{\text{oct}}$ , of the target lightcone.

evolution. These observations are contaminated by foregrounds that are several orders of magnitude stronger than the EoR signal, which the instruments are designed to detect. Recently developed component separation methods rely on maximum entropy generative models based on scattering transforms. In this work, we explored the feasibility of generative models of EoR lightcones built from 3D scattering transforms.

We constructed the 3D scattering transforms from a wavelet set obtained by taking the tensor product between a 1D wavelet set defined in  $k_z$  and a 2D wavelet set defined in the  $k_x-k_y$  plane. We demonstrated the performance of these generative models by applying them to a single simulated EoR lightcone from the LoReLi II database. We validated the generative model by comparing several realisations of the synthesised lightcone with the original target lightcone, using independent statistical estimators, including histograms, 3D and 2D power spectra, cylindrically averaged power spectra, and Minkowski functionals.

The results demonstrate that the synthesised lightcones accurately reproduce the key non-Gaussian properties and morphology of the target lightcone, including its characteristic skewed histogram, redshift-dependent structures, and morphological complexity, as probed by the Minkowski functionals. These statistics were well reproduced within the expected sample variance estimated from the data.

Beyond the overall validation, several limitations were identified:

- The synthesised lightcones reproduce the target power spectra over a wide range of scales, but clear biases appear at high- $k$  (small scales), particularly along the redshift axis. This indicates limitations of the generative model when combined with the inverse quantile transformation. However, this could be mitigated by additional post-inverse corrections or by refining the treatment of scales not constrained by the scattering transforms.
- The synthesised lightcone Minkowski functionals do not fully reproduce the  $M_2$  of the target lightcone. This suggests that the generative model struggles to produce sharp, high-contrast circular features in the target lightcone, such as ionised regions. More sophisticated scattering transform representation, such as complete scattering spectra statistics (Cheng et al. 2023), could address this, but at an additional computational cost.
- The synthesised lightcones display diverse realisations, but the variance between them may be slightly overestimated compared to the variance estimated over the octants of the target lightcone. More simulated data would be required to adequately estimate this lightcone-to-lightcone variability.

Overall, this study provides the first proof-of-concept for applying 3D scattering transform-based generative models to EoR lightcones. These results pave the way for the development of a statistical component separation method using 3D scattering transform statistics to separate the EoR signal from a realistic forward-modelled lightcone with foregrounds and noise.

Although our tests focus on EoR lightcones with a modest redshift extent, the approach is generally applicable to other spectroscopic datasets and 3D fields in which the third axis represents a different quantity. A natural next step is to apply the method to full lightcones with stronger redshift evolution across the volume.

*Acknowledgements.* The post-doctoral contract of Ian Hothi was funded by Sorbonne Université in the framework of the Initiative Physique des Infinis (IDEX SUPER) and through the overheads of Edith Falgarone’s MIST ERC. We would like to thank Florent Mertens and Adélie Gorce for their useful discussions and input. The Loreli project was provided with computer and storage resources by GENCI at TGCC thanks to the grants 2022-A0130413759 and 2023-A0150413759 on the supercomputer Joliot Curie’s ROME partition. This work was granted access to the HPC resources of MesoPSL financed by the Region Ile de France and the project Equip@Meso (reference ANR-10-EQPX-29-01) of the programme Investissements d’Avenir supervised by the Agence Nationale pour la Recherche. This research made use of astropy, a community-developed core Python package for astronomy (Astropy Collaboration 2018); scipy, a Python-based ecosystem of open-source software for mathematics, science, and engineering (Virtanen et al. 2020) – including numpy (Walt et al. 2011); matplotlib, a Python library for publication quality graphics (Hunter 2007).

## References

- Allys, E., Levrier, F., Zhang, S., et al. 2019, *A&A*, **629**, A115  
 Allys, E., Marchand, T., Cardoso, J. F., et al. 2020, *Phys. Rev. D*, **102**, 103506  
 Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123  
 Auclair, C., Allys, E., Boulanger, F., et al. 2024, *A&A*, **681**, A1  
 Becker, G. D., Bolton, J. S., Madau, P., et al. 2015, *MNRAS*, **447**, 3402  
 Blamart, M., & Liu, A. 2025, ArXiv e-prints [arXiv:2505.09674]  
 Boelens, A. M., & Tchelepi, H. A. 2021, *SoftwareX*, **16**, 100823  
 Bosman, S. E. I., Fan, X., Jiang, L., et al. 2018, *MNRAS*, **479**, 1055  
 Bruna, J., & Mallat, S. 2012, ArXiv e-prints [arXiv:1203.1513]  
 Bruna, J., & Mallat, S. 2018, ArXiv e-prints [arXiv:1801.02013]  
 Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, *SIAM J. Sci. Comput.*, **16**, 1190  
 Campeti, P., Delouis, J.-M., Pagano, L., et al. 2025, *A&A*, **700**, A136  
 Cheng, S., Ting, Y.-S., Ménard, B., & Bruna, J. 2020, *MNRAS*, **499**, 5902  
 Cheng, S., Morel, R., Allys, E., Ménard, B., & Mallat, S. 2023, ArXiv e-prints [arXiv:2306.17210]  
 Datta, A., Bowman, J. D., & Carilli, C. L. 2010, *ApJ*, **724**, 526

- Delouis, J. M., Allys, E., Gauvrit, E., & Boulanger, F. 2022, *A&A*, **668**, A122
- Eickenberg, M., Allys, E., Moradinezhad Dizgah, A., et al. 2022, ArXiv e-prints [arXiv:2204.07646]
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, **132**, 117
- Gorce, A., Douspis, M., Aghanim, N., & Langer, M. 2018, *A&A*, **616**, A113
- Gorce, A., Douspis, M., & Salvati, L. 2022, *A&A*, **662**, A122
- Greig, B., Ting, Y.-S., & Kaurov, A. A. 2022, *MNRAS*, **513**, 1719
- Hothi, I., Chapman, E., Pritchard, J. R., et al. 2021, *MNRAS*, **500**, 2264
- Hothi, I., Allys, E., Semelin, B., & Boulanger, F. 2024, *A&A*, **686**, A212
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, *MNRAS*, **389**, 1319
- Kumar, N. A., Çalıřkan, M., Hotinli, S. C., et al. 2025, ArXiv e-prints [arXiv:2506.11188]
- Lei, M., & Clark, S. 2023, *ApJ*, **947**, 74
- Levrier, F., Neveu, J., Falgarone, E., et al. 2018, *A&A*, **614**, A124
- Liu, A., Parsons, A. R., & Trott, C. M. 2014, *Phys. Rev. D*, **90**, 023018
- MacCrann, N., Qu, F. J., Namikawa, T., et al. 2024, *MNRAS*, **532**, 4247
- Mallat, S. 2011, ArXiv e-prints [arXiv:1101.2286]
- Mark Beasley, T., & Erickson, S. 2009, *Behav. Genet.*, **39**, 580
- McCullagh, P., & Tresoldi, M. F. 2020, ArXiv e-prints [arXiv:2001.03709]
- Mecke, K. R., Buchert, T., & Wagner, H. 1994, *A&A*, **288**, 697
- Meriot, R., & Semelin, B. 2024, *A&A*, **683**, A24
- Meriot, R., Semelin, B., & Cornu, D. 2025, *A&A*, **698**, A80
- Miville-Deschênes, M.-A., Lagache, G., Boulanger, F., & Puget, J.-L. 2007, *A&A*, **469**, 595
- Mousset, L., Allys, E., Price, M. A., et al. 2024, *A&A*, **691**, A269
- Munshi, S., Mertens, F. G., Koopmans, L. V. E., et al. 2025, *A&A*, **693**, A276
- Murray, S. G., & Trott, C. M. 2018, *ApJ*, **869**, 25
- Planck Collaboration XLVII. 2016, *A&A*, **596**, A108
- Prelogović, D., & Mesinger, A. 2024, *A&A*, **688**, A199
- Qin, Y., Mesinger, A., Bosman, S. E. I., & Viel, M. 2021, *MNRAS*, **506**, 2390
- Regaldo-Saint Blancard, B., Levrier, F., Allys, E., Bellomi, E., & Boulanger, F. 2020, *A&A*, **642**, A217
- Regaldo-Saint Blancard, B., Allys, E., Boulanger, F., Levrier, F., & Jeffrey, N. 2021, *A&A*, **649**, L18
- Régaldo-Saint Blancard, B., Allys, E., Auclair, C., et al. 2023, *ApJ*, **943**, 9
- Régaldo-Saint Blancard, B., Hahn, C., Ho, S., et al. 2024, *Phys. Rev. D*, **109**, 083535
- Reichardt, C. L., Patil, S., Ade, P. A. R., et al. 2021, *ApJ*, **908**, 199
- Schmalzing, J., & Buchert, T. 1997, *ApJ*, **482**, L1
- Semelin, B., Combes, F., & Baek, S. 2007, *A&A*, **474**, 365
- Semelin, B., Eames, E., Bolgar, F., & Caillat, M. 2017, *MNRAS*, **472**, 4508
- Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. 1999, *A&A*, **345**, 380
- Shimabukuro, H., Yoshiura, S., Takahashi, K., Yokoyama, S., & Ichiki, K. 2017, *MNRAS*, **468**, 1542
- Shimabukuro, H., Xu, Y., & Shao, Y. 2025, *Phys. Rev. D*, **112**, 063557
- Tiwari, H., Shaw, A. K., Majumdar, S., Kamran, M., & Choudhury, M. 2022, *JCAP*, **2022**, 045
- Valogiannis, G., & Dvorkin, C. 2022, *Phys. Rev. D*, **105**, 103534
- Vedantham, H., Udaya Shankar, N., & Subrahmanyan, R. 2012, *ApJ*, **745**, 176
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Walt, S. V. D., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22
- Watkinson, C. A., Greig, B., & Mesinger, A. 2022, *MNRAS*, **510**, 3838
- Zhang, S., & Mallat, S. 2019, ArXiv e-prints [arXiv:1911.10017]
- Zhao, X., Ting, Y.-S., Diao, K., & Mao, Y. 2023, *MNRAS*, **526**, 1699
- Zhao, X., Mao, Y., Zuo, S., & Wandelt, B. D. 2024, *ApJ*, **973**, 41

## Appendix A: Power spectrum binning

The dimensionless power spectrum of the 21cm brightness temperature  $I$  is defined as

$$\Delta_i^2(k) = \frac{A_k}{2\pi_i^2} \int |\tilde{I}(\mathbf{k}) \cdot W_i(\mathbf{k})|^2 d\mathbf{k}, \quad (\text{A.1})$$

where  $\mathbf{k}$  ( $= k_x, k_y, k_z$ ) is the wavenumber and  $W_i(\mathbf{k})$  is the spectral window function that defines the  $i$ th bin. We choose the spectral window function to be a top-hat function defined as,

$$W_i(\mathbf{k}) = \begin{cases} 1, & \text{if } \mathbf{k}_i \leq \mathbf{k} < \mathbf{k}_{i+1} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A.2})$$

where  $\mathbf{k}_i$  defines the lower wavenumber limit of the  $i^{\text{th}}$  bin. These window functions are assumed to be isotropic, so  $W(\mathbf{k}) = W(k)$  with  $k = |\mathbf{k}|$ . For the 3D power spectrum  $A_k = k^3$  and for the 2D power spectrum  $A_k = k^2$ .

Before estimating the power spectrum, it is essential to account for the non-periodic boundary conditions of the target lightcone. We address this by applying a Blackman-Harris window function to the data prior to the Fourier transform. Specifically, for the 3D power spectrum, we apply the 1D Blackman-Harris window function separately along each axis of the 3D lightcone. Similarly, for the 2D power spectrum, this window is applied independently along each axis of the respective 2D field. This windowing procedure is applied consistently to both the target  $d_i$  and synthesised  $\tilde{d}_i$  lightcones; although the  $\tilde{d}_i$  have periodic boundary conditions, this application ensures consistent treatment of boundary effects.

In addition, we also restrict the power spectra estimation to the range of scales that are well constrained by the generative model, so that the range of  $k$  modes that are binned is limited to  $\kappa_{low} \leq k \leq \kappa_{high}$  (see Section 4.3.2 for definitions).

The dimensionless cylindrically averaged power spectrum, which bins the  $k_{\parallel}$  and  $k_{\perp}$  domains separately, is defined as

$$\Delta_i^2(k_{\parallel}, k_{\perp}) = \frac{A_k}{2\pi_i^2} \int |\tilde{I}(\mathbf{k}_{\parallel}, \mathbf{k}_{\perp}) \cdot W_i(\mathbf{k}_{\parallel}, \mathbf{k}_{\perp})|^2 d\mathbf{k}_{\parallel} d\mathbf{k}_{\perp}, \quad (\text{A.3})$$

where  $A_k = k^3$ , for  $k = \sqrt{\mathbf{k}_{\parallel}^2 + \mathbf{k}_{\perp}^2}$ . We again assume an isotropic top-hat function,  $W_i(k_{\parallel}, k_{\perp})$ , for  $k_{\parallel} = |\mathbf{k}_{\parallel}|$  and  $k_{\perp} = |\mathbf{k}_{\perp}|$ , yielding

$$W_i(k_{\parallel}, k_{\perp}) = \begin{cases} 1, & \text{if } k_{\parallel,i} \leq k_{\parallel} < k_{\parallel,i+1} \ \& \ k_{\perp,i} \leq k_{\perp} < k_{\perp,i+1} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A.4})$$

where  $k_{\parallel,i}$  defines the lower wavenumber limit of the  $i^{\text{th}}$  along the  $k_{\parallel}$  domain, and  $k_{\perp,i}$  defines the lower wavenumber limit of the  $i^{\text{th}}$  along the  $k_{\perp}$  domain.