



HAL
open science

Retrieving the 21-cm signal from the Epoch of Reionization with learnt Gaussian process kernels

Florent G Mertens, Jérôme Bobin, Isabella P Carucci

► To cite this version:

Florent G Mertens, Jérôme Bobin, Isabella P Carucci. Retrieving the 21-cm signal from the Epoch of Reionization with learnt Gaussian process kernels. *Monthly Notices of the Royal Astronomical Society*, 2023, 527 (2), pp.3517-3531. 10.1093/mnras/stad3430 . hal-04179245

HAL Id: hal-04179245

<https://hal.science/hal-04179245v1>

Submitted on 12 Feb 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Retrieving the 21-cm signal from the Epoch of Reionization with learnt Gaussian process kernels

Florent G. Mertens^{1,2★}, Jérôme Bobin³ and Isabella P. Carucci^{4,5}

¹LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Université, F-75014 Paris, France

²Kapteyn Astronomical Institute, University of Groningen, PO Box 800, NL-9700AV Groningen, The Netherlands

³IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

⁴INAF-Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, I-34131 Trieste, Italy

⁵IFPU – Institute for Fundamental Physics of the Universe, Via Beirut 2, I-34151 Trieste, Italy

Accepted 2023 October 29. Received 2023 October 10; in original form 2023 July 25

ABSTRACT

Direct detection of the Cosmic Dawn and Epoch of Reionization via the redshifted 21-cm line of neutral Hydrogen will have unprecedented implications for studying structure formation in the early Universe. This exciting goal is challenged by the difficulty of extracting the faint 21-cm signal buried beneath bright astrophysical foregrounds and contaminated by numerous systematics. Here, we focus on improving the Gaussian Process Regression (GPR) signal separation method originally developed for LOFAR observations. We address a key limitation of the current approach by incorporating covariance prior models learnt from 21-cm signal simulations using variational autoencoder (VAE) and interpolatory autoencoder (IAE). Extensive tests are conducted to evaluate GPR, VAE–GPR, and IAE–GPR in different scenarios. Our findings reveal that the new method outperforms standard GPR in component separation tasks. Moreover, the improved method demonstrates robustness when applied to signals not represented in the training set. It also presents a certain degree of resilience to data systematics, highlighting its ability to effectively mitigate their impact on the signal recovery process. However, our findings also underscore the importance of accurately characterizing and understanding these systematics to achieve successful detection. Our generative approaches provide good results even with limited training data, offering a valuable advantage when a large training set is not feasible. Comparing the two algorithms, IAE–GPR shows slightly higher fidelity in recovering power spectra compared to VAE–GPR. These advancements highlight the strength of generative approaches and optimize the analysis techniques for future 21-cm signal detection at high redshifts.

Key words: methods: data analysis – methods: statistical – techniques: interferometric – cosmology: observations – reionization, first stars.

1 INTRODUCTION

Observation of the redshifted 21-cm transition line of neutral hydrogen from the Cosmic Dawn (CD; $15 < z < 30$) and the Epoch of Reionization (EoR; $z \sim 6–15$) promises to shed new light on the first billion years of our Universe. It will explore the properties of the first stellar population, the detailed timing of the heating and reionization era, as well as the sources that drove them (see e.g. Pritchard & Loeb 2012; Furlanetto 2016, for extensive reviews).

Many observational programs are underway to detect the 21-cm signal. Current radio interferometers (LOFAR,¹ MWA,² NenuFAR,³ and HERA⁴) will provide a statistical characterization of this signal

and are paving the way to the next-generation SKAO⁵ that promises an exceptional gain in sensitivity, allowing one to obtain images of these distant epochs, and, in particular, of the ionized hydrogen bubbles, sources of incomparable information on the physical conditions governing the EoR (Koopmans et al. 2015).

Detecting the 21-cm signal is extremely challenging, mainly due to the difficulty of extracting this very faint signal buried beneath astrophysical foregrounds several orders of magnitude brighter and contaminated by numerous systematics. Although the emission mechanisms of the foregrounds are well known, primarily synchrotron and free–free emission, the intrinsic chromatic response of the instrument and calibration errors create an additional impediment, introducing rapid spectral variations to otherwise spectrally regular foregrounds (e.g. Morales et al. 2012; Thyagarajan et al. 2015; Ewall-Wice et al. 2017). Mitigating these sources of chromaticity is complex. Many experiments have adopted a so-called ‘foreground avoidance’ strategy, which only looks for the signal inside a region in k -space

* E-mail: florent.mertens@gmail.com

¹Low-frequency array, <http://www.lofar.org>

²MWA, <http://www.mwatelescope.org>

³New extension in Nancy upgrading LOFAR, <https://nenufar.obs-nancay.fr/en/homepage-en/>

⁴Hydrogen Epoch of Reionization array, <http://reionization.org>

⁵SKA Observatory, <https://www.skao.int/>

where the thermal noise and 21-cm signals dominate (e.g. Parsons et al. 2012; Jacobs et al. 2016; Kolopanis et al. 2019), discarding the so-called foregrounds-wedge part where the foregrounds dominate, but at the cost of a reduced sensitivity (a factor of ~ 3 for HERA, and even a factor of ~ 6 for LOFAR). In practice, also, leakage above the wedge is often observed (e.g. Patil et al. 2016; Kern et al. 2019; Offringa, Mertens & Koopmans 2019). We refer to the recent review by Liu & Shaw (2020), and reference therein for a more detailed description of the challenge of extracting the 21-cm signal.

To take full advantage of the sensitivity of the instruments, foreground removal algorithms have been developed, particularly algorithms that do not depend on a precise foreground model, which is often difficult to obtain. For instance, blind source separation (BSS) methods have been applied in the EoR context (e.g. Chapman et al. 2012, 2013; Bonaldi & Brown 2015; Patil et al. 2017). BSS methods linearly decompose data in a set of components in pixel/visibility space modulated in frequency, only assuming statistical properties of the components (e.g. independence or sparsity). For a recent review and comparison of the most used BSS methods, see Spinelli et al. (2021).

Mertens, Ghosh & Koopmans (2018) introduced a method based on Gaussian Process Regression (GPR) that has been employed in analysing observations of various 21-cm experiments (Gehlot et al. 2019, 2020; Ghosh et al. 2020; Mertens et al. 2020) and is the method of choice of the LOFAR-EoR and NenuFAR CD projects. This approach involves constructing a statistical model prior encompassing all components contributing to the observed signal, including foregrounds, mode mixing components, 21-cm cosmological signal, and noise. The method exploits the distinct spectral signatures of these components by utilizing parameterized frequency–frequency covariance functions. Initially, the GPR method employed generic analytical covariance models commonly used in Gaussian Processes (GP) literature. However, with the increasing integration time of LOFAR observations and the prospect of an SKAO CD/EoR survey, the need for improved accuracy has driven the development of adapted and more realistic covariance models. Moreover, an excessive mismatch of the prior covariance to the true covariance may also bias the recovered 21-cm signal (Kern & Liu 2021). Here, lies the motivation and objective of this work.

Analytical formulations of foreground covariance have been developed (Trott et al. 2016; Murray, Trott & Jordan 2017). However, they are incomplete, not including, for example, the anisotropy of the galactic emission or the full complexity of the instrument’s primary beam. Concerning the 21-cm signal, analytical simulations are far too elementary, and numerical or semi-analytical codes simulating this signal are often incompatible with each other in their parametrizations. Cheng et al. (2018) have also illustrated the risks associated with employing empirically estimated covariance. This approach results in significant signal loss when utilized within the optimal quadratic estimator formalism with PAPER data. Simulations would be more suitable to provide a more physically motivated covariance, whether for the foregrounds, the systematics or the 21-cm signal. The task at hand is thus to transform a set of simulations into a covariance function, able to reproduce the covariance matrices of these simulations and replace the generic analytical covariance models used until now in our GPR formalism. Here, we consider two machine-learning (ML) methods for doing so, variational autoencoders (VAE, Kingma & Welling 2013) and Interpolatory Autoencoders (IAE, Bobin et al. 2023). We train the two with the same set of simulations and compare their performances. Finally, we test both against the Mertens, Ghosh & Koopmans (2018) standard GPR pipeline.

The paper is structured as follows. Section 2 reviews the GPR framework and introduces a generalization to learnt kernels based on VAE and IAE. Section 3 describes the 21-cm and foregrounds simulated data sets used to test the new algorithm. Section 4 reports on the results of these simulations, emphasizing the comparisons between autoencoder-based and standard GPR in various scenarios with increasing complexity.

2 METHODS

After reviewing the GPR framework, this section describes two different learning-based algorithms to build a covariance model of the 21-cm cosmological signal using simulations: VAE and IAE, Variational and IAE. The combination of the GPR algorithm to perform signal separation with ML-trained kernel is later called ML–GPR.

All methods in this work operate on a gridded visibility cube $V(\mathbf{u}, \nu)$, which is the Fourier transform along the angular direction of the observed image cube, and where $\mathbf{u} = (u, v)$ is the vector representing the coordinates in wavelength in the visibility plane and ν is the observing frequency. This choice aligns with the nature of radio interferometric observations, where visibilities are directly measured.

2.1 Gaussian process regression

In the context of 21-cm signal detection, we are interested in modelling our data d observed at frequencies x by a foreground, a 21-cm and a noise component n (Mertens, Ghosh & Koopmans 2018):

$$\mathbf{d} = f_{\text{fg}}(\mathbf{x}) + f_{21}(\mathbf{x}) + n. \quad (1)$$

The foregrounds and 21-cm signal components are expected to have distinct frequency–frequency covariance; we can then build a covariance prior model of the data composed of a foreground covariance function (also called kernel) K_{fg} , a 21-cm signal covariance function K_{21} , and a noise covariance $K_n = \text{diag}(\sigma_n^2(\mathbf{x}))$. The joint probability density distribution of the observations d and the function values f_{21} of the 21-cm signal component at the same frequencies x are then given by (Rasmussen & Williams 2005),

$$\begin{bmatrix} \mathbf{d} \\ \mathbf{f}_{21} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K_{\text{fg}} + K_{21} + K_n & K_{21} \\ K_{21} & K_{21} \end{bmatrix} \right) \quad (2)$$

using the shorthand $K \equiv K(x, x')$. The recovered 21-cm signal is then a Gaussian process, conditional on the data:

$$f_{21} \sim \mathcal{N}(\mathcal{E}(f_{21}), \text{cov}(f_{21})) \quad (3)$$

with expectation value and covariance defined by

$$\mathcal{E}(f_{21}) = K_{21} [K_{\text{fg}} + K_{21} + K_n]^{-1} \mathbf{d}, \quad (4)$$

$$\text{cov}(f_{21}) = K_{21} - K_{21} [K_{\text{fg}} + K_{21} + K_n]^{-1} K_{21}. \quad (5)$$

One may also decide to recover the fitted foregrounds component instead, $f_{\text{fg}}(x)$, and subtract it from the input data to obtain the residual, as is done in Mertens, Ghosh & Koopmans (2018) and Mertens et al. (2020).

In GPR, we perform prior covariance model selection under a Bayesian framework by choosing the model that maximizes the marginal likelihood, i.e. the integral of the likelihood times the prior. The functional form of the covariance is first selected, guided by our knowledge of the fluctuations we aim to model. We then use standard

optimization or Markov chain Monte Carlo (MCMC) methods to determine the optimal hyperparameters of the chosen covariance functions, like the coherence-scale or the variance. The marginal-likelihood is given by

$$p(\mathbf{d}|\mathbf{x}, \theta) = \int p(\mathbf{d}|\mathbf{f}, \mathbf{x}, \theta)p(\mathbf{f}|\mathbf{x}, \theta)d\mathbf{f}, \quad (6)$$

with θ being the hyperparameters of the covariance function. Under the assumption of Gaussianity, we can integrate over \mathbf{f} analytically, yielding the log marginal likelihood (LML),

$$\log p(\mathbf{d}|\mathbf{x}, \theta) = -\frac{1}{2}\mathbf{d}^\top(K + \sigma_n^2 I)^{-1}\mathbf{d} - \frac{1}{2}\log|K + \sigma_n^2 I| - \frac{n}{2}\log 2\pi \quad (7)$$

with n the number of sampled points and $K = K_{\text{fg}} + K_{21}$. The posterior probability density of the hyperparameters is then found by applying Bayes' theorem, incorporating here the prior on the hyperparameters

$$\log p(\theta|\mathbf{d}, \mathbf{x}) \propto \log p(\mathbf{d}|\mathbf{x}, \theta) + \log p(\theta). \quad (8)$$

The Matérn class of covariance functions are commonly used as prior covariance for the different components of the data. It is defined by

$$\kappa_{\text{Matern}}(r) = \frac{2^{1-\eta}}{\Gamma(\eta)} \left(\frac{\sqrt{2\eta}r}{l} \right)^\eta K_\eta \left(\frac{\sqrt{2\eta}r}{l} \right), \quad (9)$$

where K_η is the modified Bessel function of the second kind. Functions obtained with this kernel class are at least η -time differentiable. Spectrally smooth components (typically foregrounds) are best modelled with higher η , while for the 21-cm signal, we set $\eta = 1/2$ or $3/2$ (Mertens, Ghosh & Koopmans 2018). Component separation is most effective when the coherence-scale hyperparameter l of the foregrounds covariance is significantly larger than that of the 21-cm signal covariance. However, generic covariance functions may not always accurately represent the true covariance and in the low signal-to-noise regime, an inaccurate prior covariance can introduce biases in the results (Kern & Liu 2021). To address this, we can enhance the prior covariance model by learning from simulations, allowing for a more accurate representation.

2.2 Learning covariance prior model from simulations

The learning-based approach aims to build a surrogate model for the covariance that can be deployed efficiently at inference time. Covariances are derived from simulations that are fully characterized by a few parameters, for example, 3 for the 21cmFAST of the present paper. Building a surrogate model then results in computing a low-dimensional representation of the covariances. To that end, a straightforward approach would be using principal component analysis (PCA); in our tests, it did not provide satisfactory results. The latter is due to the non-linear dependency of the covariances to the underlying physical parameters. While PCA provides a low-dimensional linear representation, an accurate surrogate model mandates non-linear models. To that end, we propose investigating the application of autoencoders (AE).

We assume we have access to T simulations of the 21-cm cosmological signal that allows building a training set of covariances: $\mathcal{T}_{\text{train}} = \{K^t\}_{t=1, \dots, T}$.⁶ The prime goal of learning is to derive a low-dimensional representation or model of these covariances. From the

⁶We note here that we are not describing a full covariance, but a partial covariance averaged over some dimensions. Here, we use the frequency-frequency covariance.

mathematical modelling point of view, physically relevant covariances generally evolve smoothly, meaning they are likely to belong to a *low-dimensional manifold*. When derived from parametric models, the dimensionality of this manifold is exactly the number of free parameters in the model. However, we notice that the dependency between the underlying parameters and the measured covariances is generally *non-linear* and either unknown or challenging to handle in a component separation process. Consequently, one has to resort to learning a non-linear model of the covariances to build an efficient and reliable model, relying on an unsupervised manner.

2.2.1 Variational autoencoders

Among ML strategies, autoencoders (AE, Hinton & Salakhutdinov 2006) represent the most straightforward dimensionality reduction technique. In our context, the VAE (Kingma & Welling 2013) is particularly well adapted since it allows building a potentially complex statistical description of the data.

More formally, the set of covariances K is statistically characterized by some unknown probability distribution $p_\psi(K)$, which depends on some parameters ψ . This distribution can be described as the marginalized distribution of the joint distribution $p_\psi(K, z)$ with respect to some latent parameters z . In practice, the latent parameters will represent a low-dimensional description of the covariances. The marginalized distribution is obtained as follows:

$$p_\psi(K) = \int_z p_\psi(K|z)p_\psi(z)dz, \quad (10)$$

where $p_\psi(z)$ is the prior distribution of the latent variable and $p_\psi(K|z)$ is the conditional distribution of the covariances with respect to z . This equation defines how the distribution of K can be derived from the latent variables, eventually defining a generative model for the covariances.

In practice, the training stage corresponds to learning the parameters ψ that define the above distributions. To do so, we need a description of the conditional distribution of the latent variables with respect to K , which requires the knowledge of $p_\theta(z|K)$. However, this quantity is generally unknown and cumbersome to compute. In the framework of the VAE, it is approximated by a surrogate distribution $q_\phi(z|K)$, where ϕ stands for some parameters to be learnt.

The training stage then boils down to learning the parameters ψ and ϕ so that $q_\phi(z|K)$ is as close as possible to $p_\theta(z|K)$ according to the Kullback–Leibler divergence \mathcal{D}_{KL} . The learning loss eventually writes as (Kingma & Ba 2014)

$$\min_{\phi, \psi} -\mathbb{E}_{q_\phi} \{\log(p_\psi(K|z))\} + \mathcal{D}_{KL}(q_\phi(z|K), p_\psi(z)) \quad (11)$$

Essentially, the VAE can be described with the following ingredients:

(i) The encoder: The distribution $q_\phi(z|K)$ defines a statistical encoder that samples the latent space conditionally to the covariances. The encoder builds upon a neural network Φ to describe $q_\phi(z|K)$, and the parameters ϕ stand for the network parameters. The latent variable is then defined as follows:

$$\forall t = 1, \dots, T; \quad z_t = \Phi(K^t)$$

(ii) The decoder: As stated earlier, the distribution of the covariances $p_\psi(K)$ can be defined as the marginalization with respect to z of the joint distribution $p_\psi(K, z) = p_\psi(K|z)p_\psi(z)$. The prior distribution is generally chosen as a normal distribution, whose mean and covariance matrix are trained. Sampling the distribution $p_\psi(K)$ is then obtained by first drawing from the prior distribution of the

latent variable: $\tilde{z}^t \sim p_\psi(z)$ and then applying a non-linear function Ψ to map it back to the input domain:

$$\forall t = 1, \dots, T; \quad \hat{K}^t = \Psi(\tilde{z}^t).$$

The so-called decoder Φ is generally described with a neural network, whose parameters ϕ are learnt during training.

Details of the VAE implementation we used in the experiments are in Appendix A.

2.2.2 Interpolatory AE

As mentioned above, AE are powerful architectures to build low-dimensional non-linear signal approximations. However, their training can be difficult, especially when the training samples are scarce. A way forward to face these limitations is to impose additional constraints and better *organise* the latent space. In this context, the IAE (Bobin et al. 2023)⁷ aims to capture the geometrical structures of the space in which the data live, which is generally well approximated by some (unknown) low-dimensional manifold \mathcal{M} .

In contrast to standard AE architectures, IAE maps its latent space with so-called anchor points, which are fixed signal examples known to be on (or close to) the manifold \mathcal{M} . These anchor points are generally chosen within the training set. They serve as fixed points from which new samples can be built through a non-linear interpolation scheme.

While standard AE looks for a global description of the manifold \mathcal{M} , the IAE learns to travel on the manifold \mathcal{M} by interpolating the anchor points.

Like the VAE, the IAE architecture is defined with an encoder/decoder scheme. However, a key difference is that IAE decodes low-dimensional linear approximations of the input data in the latent space. The IAE is therefore trained so that the code of each sample can be expressed as a linear combination, or *barycenter*, of anchor points in the latent domain. Bobin et al. (2023) emphasize how this procedure allows more robustness of the model with respect to the scarcity of the training samples.

We first draw at random d samples from the training \mathcal{T} that will serve as anchor points $\{K_a^n\}_{n \in [1..d]}$. Similarly to the VAE, we define Φ and Ψ as the encoder and the decoder, respectively. The IAE networks is then defined by the three following steps:

- (i) Mapping the data to latent space: The encoder Φ maps the data to the latent space, defining some latent vector $\ell^i = \Phi(K^i)$.
- (ii) Interpolation step: Each sample is approximated as a linear interpolation from the anchor points, i.e. we compute the interpolating weights λ^i as the projection coefficients of ℓ^i on to the vector space generated by the latent codes of the anchor points $\{\Phi(K_a^n)\}_{n \in [1..d]}$:

$$\{\lambda[n]\}_{n=1, \dots, d} = \operatorname{argmin}_\lambda \left\| \Phi(K^t) - \sum_{n=1}^d \lambda[n] \Phi(K_a^n) \right\|_2^2 \quad (12)$$

The interpolated samples are then defined as

$$\tilde{\ell}^t = \sum_{n=1}^d \lambda^t[n] \Phi(K_a^n). \quad (13)$$

The interpolation step is denoted as I so that $\tilde{\ell}^t = I \circ \Phi(K^t)$.

(iii) Mapping the interpolated samples back to the data domain: The decoder ψ maps back the interpolated samples from the latent space to the data domain: $\tilde{K}^t = \Psi(\tilde{\ell}^t)$.

⁷<https://github.com/jbobin/IAE>

In this model, both the encoder and the decoder are learnt. Model optimization is done by minimizing the reconstruction error between the training data $\{K^t\}_j$ and their approximation, i.e.:

$$\operatorname{argmin}_{\phi, \psi} \sum_j \left\| K^t - \Psi \left(\sum_{n=1}^N \lambda^j[n] \Phi(K_a^n) \right) \right\|_2^2 + \beta \sum_j \left\| \Phi(K^t) - \sum_{n=1}^N \lambda^j[n] \Phi(K_a^n) \right\|_2^2. \quad (14)$$

Once the IAE is learnt, any point of \mathcal{M} can be approximated as an interpolant of the anchor points in the code domain back-projected to the sample domain:

$$\mathcal{M} \approx \left\{ x, \exists \{\lambda[n]\}_{n \in [1..N]} \in \mathbb{R}^N, x = \psi \left(\sum_{n=1}^N \lambda[n] \Phi(K_a^n) \right) \right\}. \quad (15)$$

Details of the architecture we used in the experiments as well as its implementation, are described in Appendix A.

2.2.3 Application of AE-based models for covariance estimation

AE-based models can be interpreted and applied in different ways. To that respect, it is customary to use AE-based models to compute non-linear, low-dimensional approximations of the classes of signals as sampled by the training set. For instance, once learnt, one can use the IAE model to produce a low-dimensional approximate of some signal K by composing the application of the encoder, the interpolator and the decoder:

$$\hat{K} = (\Psi \circ I \circ \Phi)(K).$$

We can interpret this operation as a projection on to the low-dimensional manifold implicitly defined by the encoder and the decoder.

From a different viewpoint, the decoder of AE-based models boils down to a learnt parametric model for the covariances. Combined with a sampling process, as in the VAE, it leads to generative models.

Let us now consider the case of observations involving the 21-cm signal along with foreground and noise components. The observation model is described as follows:

$$K_d = K_{\text{fg}} + K_{21} + K_n.$$

In the scope of covariance estimation from observations, one can substitute standard analytic kernels, such as Matérn or Gaussian kernels, with the decoder, which can be treated as any parametric model $\Psi(\lambda)$, where λ lives in the latent space of a given AE-based model. The observation model then reads as

$$K_d = K_{\text{fg}} + \Psi(\lambda) + K_n.$$

Covariance estimation then boils down to estimating the hyperparameter λ during the inference procedure. The *a posteriori* distribution now reads as

$$\log p(\theta, \lambda | \mathbf{d}, \mathbf{x}) \propto \log p(\mathbf{d} | \mathbf{x}, \theta, \lambda) + \log p(\theta) + \log p(\lambda), \quad (16)$$

where the LML is defined in equation (7), and where the hyperparameters λ are related explicitly to the parameterization of the AE-based models that act as the covariance of the 21-cm signal. The hyperparameters θ are related to modelling additional components such as noise and foregrounds, depending on the model we consider (see Section 2.1).

To compute the maximum a posteriori and estimate uncertainties on the hyperparameters θ and λ , we use a Markov chain Monte Carlo (MCMC). Such methods aim to sample the model parameters' posterior probability distribution given the observed data. In practice, this is performed with the ensemble sampler algorithm of Goodman & Weare (2010), implemented in the `emcee` python package (Foreman-Mackey et al. 2013).

2.2.4 Training the 21-cm signal kernel

In the generic description of the learnt kernel, the kernel is trained on a set of covariance K . However, when it comes to the specific case of the 21-cm signal, we can exploit the isotropic nature of the signal and train the AE on a set of 21-cm power spectra, denoted as $P_{\text{HI}}(k)$. This training is equivalent to training on the covariance K . With the AE, we learn the shape of the power spectra while keeping the variance as a separate parameter. Therefore, during the pre-processing of our learning step, our 21-cm power spectra of the training set are scaled to have the same power. The trained decoder of the AE thus takes n parameters as input; it then outputs a normalized spherically averaged power spectrum, which is then converted to frequency–frequency covariance matrix, using the expression of the angular power spectrum, in the flat-sky approximation (Datta, Choudhury & Bharadwaj 2007):

$$C_l(\Delta\nu, k_{\perp}) = \frac{T^2}{\pi r_v^2} \int_0^{\infty} P_{\text{HI}} \left(\sqrt{k_{\perp}^2 + k_{\parallel}^2} \right) \cos(k_{\parallel} r_v \Delta\nu) dk_{\parallel}, \quad (17)$$

where T is the differential brightness temperature of the 21-cm signal, r_v^2 is the comoving distance at the mean frequency of the observation cube. Additionally, k_{\perp} and k_{\parallel} represent the Fourier modes in the transverse (spatial) and line-of-sight (frequency) directions, respectively. A more comprehensive description of these concepts and their definitions can be found in Section 2.3. Combining these elements, we form the 21-cm trained kernel, which takes $n + 1$ parameters (n corresponding to the dimension of the latent space of the AE), and one parameter for the variance of the kernel. We note here the dependence of C_l on k_{\perp} which provides a significant advantage over the generic kernel traditionally used in the standard GPR method as later tests will demonstrate.

2.3 Power spectra estimation

Current generation of redshifted 21-cm experiments aims at characterizing the power spectra of the signal at different redshift to quantify the scale-dependent second moment of the signal and its evolution as function of redshift. We compute it by taking the Fourier transform of the gridded visibility cubes $V(\mathbf{u}, \nu)$ in the frequency direction. We can then define the cylindrically averaged power spectrum as (Mertens, Ghosh & Koopmans 2018):

$$P(k_{\perp}, k_{\parallel}) = \frac{X^2 Y}{\Omega_{\text{PB}} B} \left\langle |\hat{V}(\mathbf{u}, \tau)|^2 \right\rangle, \quad (18)$$

where $\hat{V}(\mathbf{u}, \tau)$ is the Fourier transform in the frequency direction, B is the frequency bandwidth, Ω_{PB} is the integral of the square of the primary beam, X and Y are conversion factors from angle and frequency to comoving distance, and $\langle \dots \rangle$ denote the averaging over baselines of the same length. The Fourier modes are in units of inverse comoving distance and are given by (Morales, Bowman & Hewitt 2006; Trott, Wayth & Tingay 2012)

$$k_{\perp} = \frac{2\pi|\mathbf{u}|}{D_M(z)}, \quad (19)$$

$$k_{\parallel} = \frac{2\pi H_0 v_{21} E(z)}{c(1+z)^2} \tau, \quad (20)$$

$$k = \sqrt{k_{\perp}^2 + k_{\parallel}^2}, \quad (21)$$

where $D_M(z)$ is the transverse co-moving distance, H_0 is the Hubble constant, v_{21} is the frequency of the hyperfine transition, and $E(z)$ is the dimensionless Hubble parameter (Hogg 1999). We also define the dimensionless power spectrum by averaging the power spectrum in spherical shells as

$$\Delta^2(k) = \frac{k^3}{2\pi^2} P(k). \quad (22)$$

The latter is well suited for characterizing the 21-cm signal, which we expect to be isotropic at a given redshift. We limit our analysis to a bandwidth of 12 MHz to limit the effects of signal evolution, which would introduce anisotropy (Mertens et al. 2020).

To compute the power spectrum and its uncertainty for a specific component, say the 21-cm component, we start by taking m samples from the posterior distribution of the hyperparameters, that we can take directly from the MCMC chain. We then generate $\mathcal{E}(f_{21})$ and $\text{cov}(f_{21})$ using equation (4) and equation (5), and produce the power spectrum of $(\mathcal{E} + \delta_{21})$ with δ_{21} a sample drawn from the Gaussian distribution with covariance $\text{cov}(f_{21})$. We then compute the power spectrum and associated 1σ uncertainty of the 21-cm component by taking the median and standard deviation of the m power spectra thus generated.

3 SIMULATED DATA

This section provides a description of the simulated data sets used for testing the new ML-GPR method. We focus on SKA-low and simulate a gridded visibility cube using the planned SKA array configuration and the expected SKA thermal noise level. Noise simulation is performed using `ps_eor`.⁸ To produce the 21-cm signal component, we use 21cmFAST (Mesinger & Furlanetto 2007; Mesinger, Furlanetto & Cen 2011) simulated cubes for both the training set of the learnt kernel and the reference signal that we will attempt to recover. To model the foreground component, we use a similar strategy as in Mertens, Ghosh & Koopmans (2018) and use GP with a frequency–frequency covariance model expected for an SKAO observation. As our primary focus is on validating the 21-cm learnt kernel, we use a simple foreground simulation generated with the same GP covariance function used for the signal separation to avoid bias from a mismatched foregrounds kernel: The same covariance model will be used as a foreground component in our GP model. However, we plan to try more realistic foreground simulations in future work.

The simulation spans a frequency range of 122–134 MHz with a spectral resolution of 0.2 MHz, corresponding to a redshift $z \sim 10.1$. The maps cover a field of view of 4.5° with a baseline ranging from 35 to 300λ , which gives us a maximum angular resolution of 10 arcmin. We chose these parameters to get an optimal sensitivity to the 21-cm signal.

3.1 Foregrounds

Following Mertens, Ghosh & Koopmans (2018), we model foregrounds using two components: the intrinsic astrophysical fore-

⁸https://gitlab.com/flomertens/ps_eor

grounds and the instrumental effects, also known as the mode-mixing component.

At the radio frequencies at which 21-cm experiments operate, the astrophysical foregrounds emission foregrounds can be classified into two primary groups: (i) galactic foregrounds, which are primarily linked to the diffuse synchrotron and, to some degree, free-free emission from the Milky Way; and (ii) extragalactic foregrounds, which are linked to radio emission from star-forming galaxies and active galactic nuclei, and to a lesser extent, radio halos and relics. The emission mechanisms are well known and are characterized by a power-law spectrum. Our approach in this work is not to model the spatial distribution of astrophysical foregrounds. Instead, we simulate this component using a Gaussian covariance with an 80-MHz frequency–frequency coherence scale. We do so because the ML–GPR method mainly accounts for frequency–frequency correlations (also we account for some spatial information through the dependence of the frequency–frequency covariance to the baseline length). Hence, our treatment is appropriate, given the limitations of our current implementation. The variance of this component is matched to the variance observed in LOFAR observations after direction-dependent calibration and sky model subtraction (see fig. 7 in Mertens et al. 2020).

We use a Gaussian Process to simulate instrumental effects, which can arise from various sources, such as instrument chromaticity and calibration imperfections. Our approach is motivated by the analysis of LOFAR data, which suggests that medium-scale fluctuations can be accurately modelled using a Matérn covariance function, as previously done in Mertens, Ghosh & Koopmans (2018). The mode-mixing typically appears in k_{\perp} , k_{\parallel} -space as a wedge-like structure (Datta, Bowman & Carilli 2010; Morales et al. 2012), with the level of spectral fluctuations increasing as the baseline length grows. To address this effect, we scale the frequency–frequency coherence scale with the angular scale, with values ranging from 4 MHz for a baseline length of 35λ to 2 MHz for a baseline length of 300λ . Here, we set $\eta = 5/2$ for the Matérn function. In our study, the primary focus is on modelling the 21-cm signal. While analytically defined covariance for instrumental effects has been developed in prior works (Liu, Parsons & Trott 2014a,b), we opted here for a simple method to generate foregrounds, including instrumental effects).

In this work, we also investigate the resilience of our method to recover the 21-cm signal in the presence of additional, less spectrally smooth components. Specifically, we consider an excess component observed in LOFAR observations (Mertens et al. 2020) that cannot be attributed to astrophysical foregrounds, instrumental effects, or the 21-cm signal. To simulate this component, we use a Matérn kernel with $\eta = 5/2$ and a coherence scale of 0.4 MHz, which captures its statistical characteristics. We aim to evaluate the robustness of our method in handling such contaminants.

3.2 21-cm signal

We simulate the 21-cm signal with 21cmFAST (Mesinger & Furlanetto 2007; Mesinger, Furlanetto & Cen 2011), a semi-numerical code that characterizes the relevant physical processes with approximate methods, resulting in faster and less expensive computations than accurate radiative transfer simulations. The semi-analytic codes generally agree well with hydrodynamical simulations for comoving scales > 1 Mpc.

Several parameterization options are available in 21cmFAST (e.g. Park et al. 2019). In our study, we opt for the simpler three-parameter parameterization, where ζ_0 represents the ionizing efficiency of high- z galaxies, T_{vir} denotes the minimum virial temperature of star-




Figure 1. Spherically averaged power spectra of 450 selected 21cmFAST simulations from the training set. The total training set is made of 5000 simulations. In red is shown the power spectra of the reference signal.

forming haloes, in units of M_{\odot} , and R_{mfp} which represents the mean free path, in cMpc, of ionizing photons within ionizing regions. We produce a total of 5000 simulation co-eval cubes at $z \sim 10.1$, with ζ_0 ranging between 10 and 250 with a step of 5, $\log(T_{\text{vir}})$ ranging between 4 and 6 with a step of 0.1 and R_{mfp} ranging between 5 and 25 cMpc with a step of 2 cMpc. For the fiducial simulation, we use the parameters $[\zeta_0, \log(T_{\text{vir}}), R_{\text{mfp}}] = [102, 4.65, 14]$, outside of the training set. In Fig. 1, we plot in red the spherically average power spectrum of the fiducial simulation against the spectra of some of the simulations from the training set (grey lines).

3.3 Noise

We test the signal separation algorithm in the context of a SKA-Low observation. With this aim, we simulate realistic SKA thermal noise using the `ps_eor` code. At completion, the SKA-Low will be made of 512 stations, each of 256 elements, providing exquisite sensitivity at our scale of interests. Our simulation process involved creating the uv -coverage of the instruments by simulating the uv -tracks of each individual baseline and then gridding them to a uv -grid. The uv -coverage grid is simulated assuming the current plan for antennae distribution of SKA-Low.⁹ We simulate the uv -tracks corresponding to a 10-h long observation, with an integration time $t_{\text{int}} = 10$ s targeting the Murchison Widefield Array (MWA) EoR-0 deep field (RA = 0h00, Dec. = -27°) (Trott et al. 2020), which is a potential future target for SKA-Low observations. The noise in each uv -cell, and for each frequency channel ν is then generated following a Gaussian distribution with standard deviation:

$$\sigma(u, v, \nu) = \frac{k_{\text{B}} T_{\text{sys}}(\nu)}{A_{\text{eff}}} \sqrt{\frac{1}{2\delta_{\nu} N(u, v, \nu) t_{\text{int}}}}. \quad (23)$$

with T_{sys} is the system temperature,¹⁰ A_{eff} is the effective collecting area, δ_{ν} is the spectral resolution (0.2 MHz in our case) and $N(u, v, \nu)$ is the number of visibilities that are fall into a uv -coverage grid. Noise cube emulating the integration of up to $t_{\text{obs}} = 1000h$,

⁹The SKA-low design is given at https://www.skao.int/sites/default/files/documents/d18-SKA-TEL-SKO-0000422.02_SKA1_LowConfigurationCoordinates-1.pdf.

¹⁰The SKA LFAA station design report is available at <https://arxiv.org/abs/2003.12744v1>

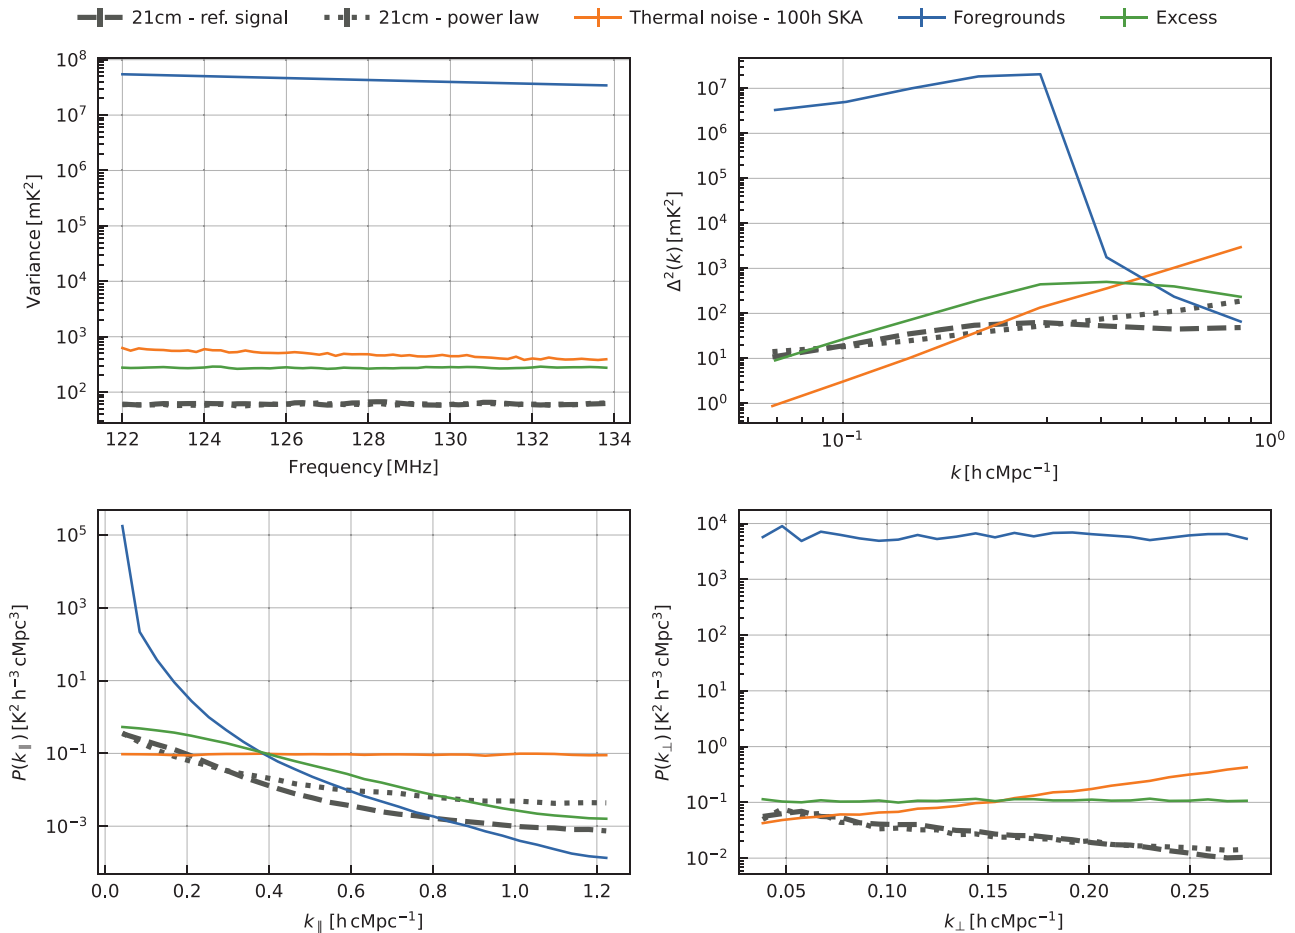


Figure 2. Power spectra of all simulated components used in our tests: foregrounds (blue), SKA thermal noise for 100h of observation (orange), excess component (green), reference 21-cm signal (dashed black line), and synthetic 21-cm signal (dotted black line). We show their variance as function of frequency (top-left panel), the spherically averaged power spectra (top right panel), the cylindrically averaged power spectra as function of $k_{||}$ (averaged over k_{\perp} , bottom left panel) and as function of k_{\perp} (averaged over $k_{||}$, bottom right panel).

of observations are generated by scaling down the noise cube by a factor of $\sqrt{(10/t_{\text{obs}})}$.

3.4 Tests scenarios

In Fig. 2, we summarize all the components of our reference simulation and their differences in terms of variance (top left panel), spherically average power spectrum (top right) and cylindrically average power spectra averaged (bottom panel). We also include a different, *synthetic*, power-law-like 21-cm signal that we will later employ in additional tests.

We assess the performance of the ML-GPR methods by testing them against various scenarios which are made of different components:

Without foregrounds: This test will evaluate the raw performance of the algorithm with an ideal scenario including only the 21-cm signal and noise components.

With foregrounds: Evaluating the impact of foregrounds by including the 21-cm signal, noise, and foregrounds.

Without 21-cm signal: This will be a null test in which no 21-cm signal is present in the data, just noise and foregrounds.

With an excess component: In addition to the foregrounds, an extra component is added to the observation. In this case, we will test two possibilities: if the excess is accounted for or not in the GPR covariance model.

Synthetic 21-cm signal: Finally, the robustness of the methods is tested in recovering a class of models which are very different from what is used in the training set by employing a power-law-like 21-cm signal.

For all these scenarios, we will compare the performance of the IAE and VAE kernels to the approach of Mertens, Ghosh & Koopmans (2018), which uses an analytically defined covariance.

4 RESULTS

In this section, we first evaluate the reconstruction performance of the VAE and IAE and we later conduct various 21-cm signal recovery tests across different scenarios.

4.1 Model evaluation

The quality of the VAE- and IAE-based models are first evaluated independently of the regression task that we will eventually do.




Figure 3. Normalized MSE distribution between the estimated and true power spectra from the test set for the VAE (orange) and IAE (blue) models, for $n = 2$ (top panel) and $n = 3$ (bottom panel).

The goal is to assess the trained models' accuracy and ability to capture the structure of the 21-cm power spectra. Whether it is IAE or VAE, both methods are used to build a low-dimensional and non-linear signal representation of the 21-cm. To that purpose, these two approaches are evaluated with respect to two key parameters:

- (i) The dimensionality n of the latent space. n sets the complexity of the trained models. An accurate low-dimensional model is important because a compressed representation allows for better discrimination between the 21-cm signal and spurious contamination (e.g. foregrounds, noise, etc.).
- (ii) The size of the training set. Evaluating the reliability of the trained model with respect to the size of the training set is crucial because the ability to train accurate models from scarce training samples will make possible the applications to more realistic EoR simulations, whose computational cost limits the number of samples for an eventual training stage.

Fig. 3 displays the normalized mean squared error (NMSE) in decibels of the estimated 21-cm signal power spectrum across the test set for latent space dimensionality $n = 2$ and 3. Comparing the top and bottom panels, and as expected, increasing the dimensionality improves the accuracy of trained models for the IAE, especially when the training set is large enough (e.g. the number of samples larger than 500 in this experiment). The improvement for the VAE is very marginal. For a small training set, the accuracy of the models is likely limited by the limited number of available training samples.

For $n = 2$ and 3, and as expected, both the IAE and VAE provide increasingly more accuracy as the number of training samples increases. The IAE provides more accurate models in all cases, with a significant NMSE¹¹ gain close to 10 dB for $n = 3$ and a number of

¹¹The NMSE of an estimate \hat{x} of a signal x is defined as $-20 \log(\|x - \hat{x}\|_2 / \|x\|_2)$.




Figure 4. Normalized fractional ratio distribution between the estimated and true power spectra from the test set for the VAE (orange) and IAE (blue) models, for $n = 2$ (top panel) and $n = 3$ (bottom panel).

training samples of 5000. The scatter of the NMSE across the test set is similar for both methods and decreases slightly when the number of training set samples increases.

The NMSE measures the discrepancy between the estimated and input power spectra, which will be dominated by errors on its largest values located at large k modes. The relative error ratio displayed in Fig. 4 is more sensitive to the lowest values of the power spectra (i.e. at small k modes). In this case, the VAE and IAE yield quite comparable average accuracy. Except for the case of 50 training samples, the IAE yields slightly smaller dispersion across the test set. Considering the limited improvement gained by the higher dimensionality of the latent space and the increased complexity in the covariance model, we will focus on the $n = 2$ case for the remainder of this paper.

In Fig. 5, we present the reconstructed power spectra corresponding to the best and worst cases of relative absolute errors for both VAE and IAE. For these experiments, we use the model trained with 5000 samples and $n = 2$. Overall, the IAE tends to produce more accurate low-dimensional models of the 21-cm power spectra than the VAE. Even when the VAE exhibits the best reconstruction performance, the IAE performs similarly. Moreover, the relative absolute error of the worst IAE case is significantly lower than that of the worst VAE case. In general, The VAE and IAE generally perform better when dealing with 'monotonous' power spectra, while their performance degrades when confronted with power spectra that contain more complex features. The next sections will focus on extending these comparisons in the context of power spectra estimation from noisy and contaminated signals.

4.2 21-cm signal power spectrum estimation

We now evaluate the performance of the VAE and IAE models in retrieving the 21-cm EoR signal from a contaminated data cube and compare them with the standard GPR method (using a Matern




Figure 5. Reconstructed spherically averaged power spectra (top panel) and residual (bottom panel) in the best VAE case (first panel), best IAE case (second panel), worst VAE case (third panel), and worst IAE case (fourth panel), estimated across the test set, showing both the VAE (blue) and IAE (orange) results for the four cases. For comparison, we show in grey the input signal. The averaged fractional absolute reconstruction error is also shown as a dashed line.

covariance for the 21-cm component with $\eta = 3/2$. We use the VAE and IAE models trained with dimensionality $n = 2$ and a training set of 5000 samples. We will refer to the ML-GPR method employing the VAE model as VAE-GPR. Similarly, the ML-GPR method using the IAE model will be IAE-GPR.

4.2.1 Scenario I: 21-cm signal and noise

We begin by considering a simple scenario involving a data cube comprising only instrumental noise and the 21-cm signal. The results of the reconstruction are presented in the left panels of Fig. 6, where we compare the spherically averaged power spectra (top panel) and the cylindrically averaged power spectra along the line of sight (middle panel) and perpendicular to it (bottom panel). Even in this simple setup, both the VAE-GPR (in blue) and IAE-GPR (in orange) methods outperform the traditional GPR (in green), effectively recovering the input 21-cm signal. The GPR approach tends to underestimate the power at larger scales while overestimating it at smaller scales, particularly evident in the perpendicular direction. Interestingly, the scale (or k_{\perp}) dependence inherent in the new methods proves advantageous, contrasting with the lack of scale dependence in standard GPR, which results in overfitting at large and underfitting at small angular scales. Moreover, we report minimal uncertainties in the reconstruction, with slightly larger uncertainties at smaller scales (larger k) due to the lower signal-to-noise ratio (SNR) in those regions. It is important to note that the uncertainty in

the 21-cm signal estimated by GPR is smaller than the thermal noise uncertainty because of the additional constraint provided by the 21-cm signal prior. The reconstructed power spectra of the 21-cm signal result from a regression process in which the prior knowledge of the 21-cm signal can influence the outcome, particularly in low SNR scenarios, such as the thermal noise-dominated part of the power spectra. In this context, our approach can be considered model-dependent. However, our framework incorporates specific design choices, such as the independence between variance and shape parameters, the diverse range in the simulation training set, and the utilization of generative neural networks. These choices enable our method to be flexible and accommodate various possibilities, as our latest scenario demonstrates.

4.2.2 Scenario II: 21-cm signal, noise and foregrounds

In the middle column of Fig. 6, we consider the scenario where the data cube comprises the 21-cm signal, noise, and foregrounds, representing the ideal case without additional systematics in the data. Despite the presence of larger uncertainties, both the VAE-GPR and IAE-GPR methods successfully recover the 21-cm signal. In contrast, the standard GPR approach exhibits a worsened deviation from the ground truth. Notably, the limitations of the standard GPR arise from its inability to accurately capture the shape of the input power spectra, primarily due to the lack of baseline dependence of the covariance prior model. The latter is particularly

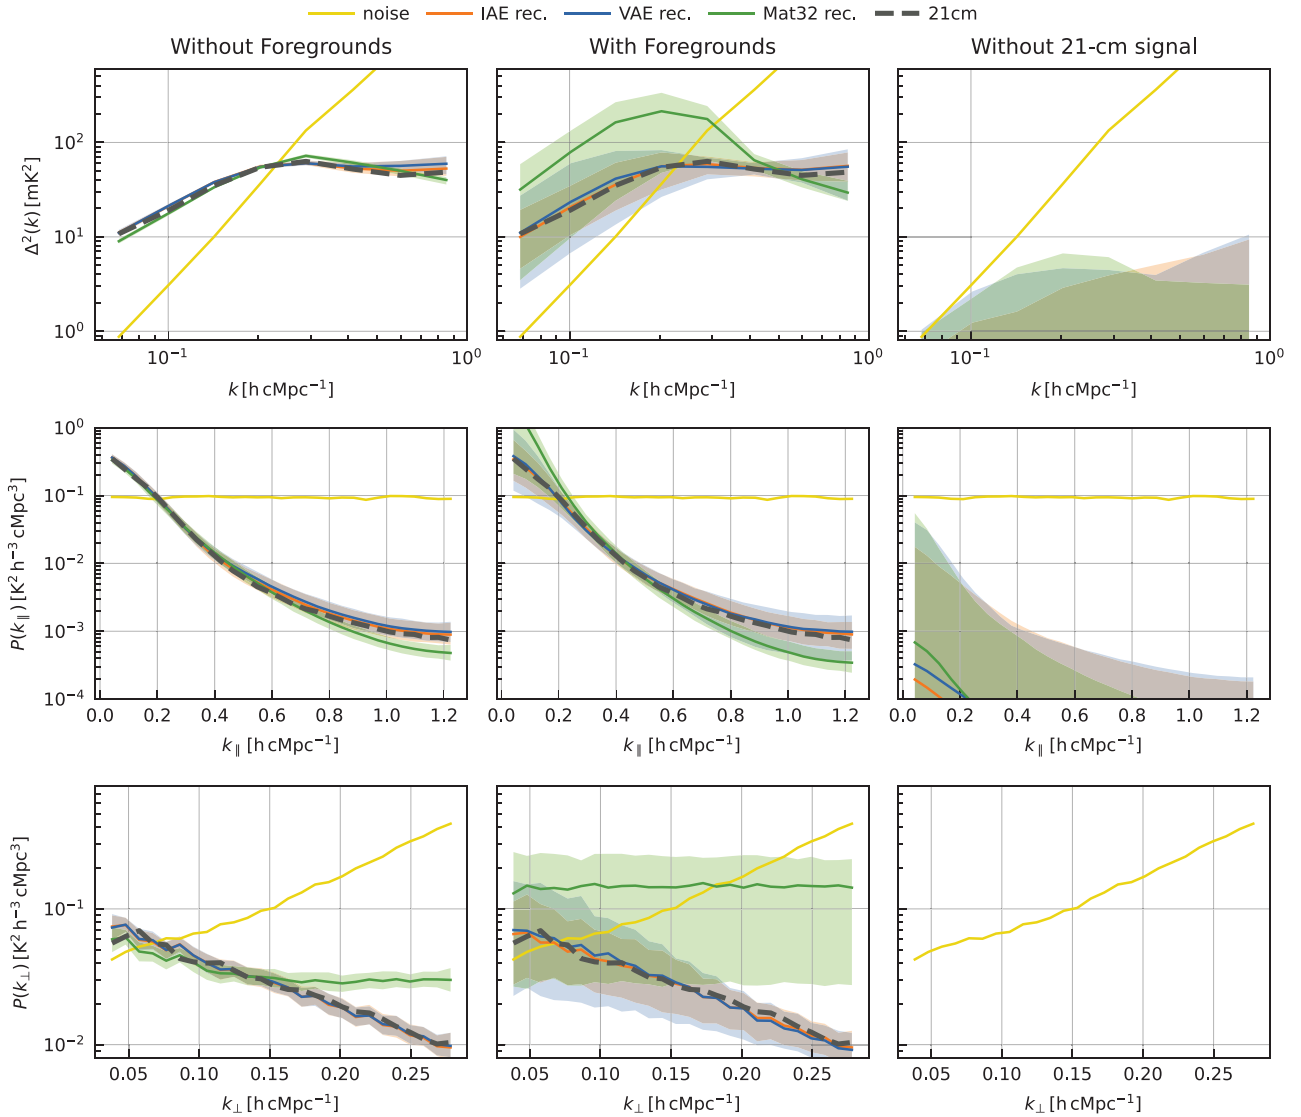


Figure 6. Recovered 21-cm signal power spectra using different methods: GPR (green), VAE–GPR (blue), and IAE–GPR (orange). The left panel shows the results when the data consist of only noise and the 21-cm signal. In the middle panel, a foregrounds component is added to the data, and in the right panel, the 21-cm component is absent. The power spectra of the noise component are shown in yellow, and the dashed black line represents the power spectra of the 21-cm component. The shaded area represents the 2σ uncertainty on the recovered 21-cm signal. The top panel displays the spherically averaged power spectra, while the center and bottom panels show the cylindrically averaged power spectra as a function of $k_{||}$ and k_{\perp} (averaged over k_{\perp} and $k_{||}$, respectively).

evident when considering the power spectrum as a function of k_{\perp} (bottom panel), which presents a flat behaviour. In contrast, the ML–GPR methods perfectly recover the shape of the input power spectra. This result can be attributed in part to the addition of the spatial information through the dependence of the frequency–frequency covariance to the base-line length. However, this is not the sole factor in the enhanced performance of the ML–GPR method. Relying solely on an Exponential or Matern covariance function with k_{\perp} dependence would not suffice. The pivotal advantage of incorporating the autoencoder in this context lies in its ability to establish a more informative prior for the 21-cm signal. Even with k_{\perp} dependence, Exponential or Matern covariance functions might not comprehensively capture all the potential 21-cm signal frequency–frequency covariance. These functions – inherently generic – might not encapsulate the intricate complexity of the intrinsic covariance in the 21-cm signal. The uncertainty is highest in regions dominated

by foregrounds ($k < 0.3 \text{ h cMpc}^{-1}$) or noise ($k > 0.5 \text{ h cMpc}^{-1}$). Comparing the performance of the IAE and VAE, both methods exhibit equivalent recovery, although the IAE demonstrates a slightly smaller bias.

4.2.3 Scenario III: noise only

In this particular scenario, we evaluate the performance of the algorithm when the 21-cm signal is significantly below the sensitivity threshold of the instrument, resulting in a data cube consisting only of foregrounds and noise components. When looking at the result of the MCMC, in all three cases (VAE–GPR, IAE–GPR, and standard GPR) the hyperparameters that control the shape of the power spectra (the two latent space parameters for the VAE and IAE, and the Matern coherence-scale for the standard GPR) are unconstrained: There




Figure 7. Recovered 21-cm signal power spectra using different methods: GPR (green), VAE-GPR (blue), and IAE-GPR (orange). In the left panel, an excess component is introduced into the data without being accounted for in the GPR model. The middle panel demonstrates the results when the excess component is properly accounted for. In the right panel, the reference 21-cm signal is replaced with a synthetic 21-cm signal modelled by power-law power spectra. The power spectra of the noise component are shown in yellow, and the dashed black line represents the power spectra of the 21-cm component. The shaded area represents the 2σ uncertainty on the recovered 21-cm signal. The top panel displays the spherically averaged power spectra, while the center and bottom panels show the cylindrically averaged power spectra as a function of $k_{||}$ and k_{\perp} (averaged over k_{\perp} and $k_{||}$, respectively).

posterior is flat. The variance of this component is also unconstrained with a large right-tail which is set by the sensitivity (all variance hyperparameters are inferred in log-space). This is the reason why, also the recovered power spectra in the left panel of Fig. 6 is below the noise power spectra (yellow line), the corresponding 2σ uncertainty is large and reach region which can be just below the noise. This is the case for all three methods.

4.2.4 Scenario IV: 21-cm signal, noise, foregrounds, and extra component

To assess the robustness of our algorithm against more realistic systematic effects, we introduce an additional foreground component into the data cube, as described in Section 4.2. This component is designed to mimic the statistical properties of the ‘excess’ observed

in current LOFAR–EoR analysis (Mertens et al. 2020). Notably, this excess component is currently not subtracted from LOFAR data, making it a limiting factor in the current analysis. While improved calibration and processing techniques may potentially mitigate or eliminate this excess component, the ability to robustly separate it from the 21-cm signal is crucial for maximizing the experiment’s sensitivity.

In the initial analysis, we examine the recovered 21-cm signal without including any additional component in our prior covariance model, considering only the foregrounds, noise, and 21-cm component. The results presented in the first column of Fig. 7 indicate poor performance for all three methods. We find that the recovered signal comprises both the excess component and the 21-cm signal. The excess present in the data is being absorbed by additional variance in the 21-cm component. Since this excess component was not

accounted for in the covariance prior model, the algorithm effectively attributes the unexplained variance to the 21-cm signal. However, we also observe a discrepancy in the shape of the recovered signal power spectra compared to the shape of the excess power spectra. This disparity arises from the fact that the covariance prior model for the 21-cm signal does not include this specific class of shape.

To address this limitation, we introduce a second step in our analysis where the same Matern covariance used to simulate the excess component is added to the GPR covariance prior model. By doing so, we account for the statistical properties of the excess component. This assumes that we have a thorough understanding of the systematic and can perfectly account for it in our analysis. The results (middle panel of Fig. 7) demonstrate that while the uncertainty in the power spectra significantly increases, the input signal lies well within the 2σ uncertainty range of the recovered signal. Both IAE-GPR and VAE-GPR show successful recovery, while standard GPR performs significantly worse. We observe a greater mixing between the excess and 21-cm components in standard GPR, resulting in a stronger bias. This discrepancy is particularly pronounced when examining the power spectra in the k_{\perp} direction, highlighting the importance of properly accounting for all components, including systematics, in the covariance prior model.

4.2.5 Scenario V: synthetic 21-cm signal, noise, and foregrounds

In the last test scenario, we aim to recover a completely synthetic and non-physical model, which represents a case where the true power spectra significantly deviate from current model predictions, and is thus not present in the training set. To simulate this scenario, we adopt a power-law model with the equation $\Delta^2(k) = \sigma_{21}^2 k^2$, where σ_{21}^2 is chosen to match the variance of the fiducial 21-cm signal. The results of this test, presented in the last column of Fig. 7, reveal that while none of the methods are able to perfectly recover the 21-cm signal, the IAE-GPR, and VAE-GPR methods exhibit better recovery of the shape compared to the standard GPR method. This improvement is particularly evident for k modes greater than 0.1, h , $cMpc^{-1}$. Upon examining the power spectra as a function of k_{\parallel} , we observe that the discrepancy between the input and recovered signal is more pronounced for modes where the foregrounds dominate, specifically the low k_{\parallel} modes.

This test underscores the importance of having a sufficiently wide prior for the 21-cm signal in order to account for unexpected 21-cm signal. It also highlights the need for a feedback loop with theoretical models, where simulations can be updated to explain observations, signal separation can be re-run with the updated models, and the resulting signal can be refined accordingly.

4.2.6 Effect of varying the noise level

In this section, we explore the influence of varying thermal noise levels on the recovery of the 21-cm power spectra. To investigate this, we conduct simulations with different integration times, ranging from 50 to 1000 h. For this analysis, we used the same noise realization, which is subsequently scaled to match the desired integration time level. This was done to prevent any fluctuations in the results solely due to different noise realizations. The results of this analysis are presented in Fig. 8, illustrating the ratio between the recovered and input 21-cm power spectra using the VAE- and IAE-trained kernel.

Our findings indicate that reducing the thermal noise level has minimal impact on the uncertainty of the recovered signal for the k -modes dominated by foregrounds ($k < 0.3 h cMpc^{-1}$). This suggests that the




Figure 8. Ratios of recovered over input 21-cm power spectra using a VAE (top panel) and IAE (middle panel) trained with 5000 samples and for different level of noise in the data: between 50 and 1000 h of equivalent SKA noise. Each case is depicted with a small offset around the measured k -bin for clarity. The 2σ uncertainty is also indicated for each case. In the bottom panel, the power spectra of the noise for the different cases are shown, along with the power spectra of the 21-cm signal, illustrating the SNR due to the noise. Foregrounds dominate the power spectra for $k < 0.3 h cMpc^{-1}$.

primary source of uncertainty in the foreground-dominated region stems from the foreground itself rather than the noise. However, for k -modes where noise dominates ($k > 0.4 h cMpc^{-1}$), increasing the integration time substantially decreases the uncertainty. These results underscore the crucial role of foregrounds in contributing to the uncertainty in the recovery of the 21-cm power spectra.

4.2.7 Varying the size of the training set

In this section, we examine the influence of the training set size on the recovery of the 21-cm signal using our trained kernel. While gathering a large training set is feasible for seminumerical simulation codes like 21cmFast, it may become computationally prohibitive for full 3D radiative transfer simulation codes (e.g. Ciardi et al. 2001; Semelin et al. 2017). However, the generative nature of our methods enables us to achieve satisfactory recovery even with a relatively small data set.

We assess the behaviour of the recovery using different training set sizes, and the results are presented in Fig. 9. Remarkably, even with only 100 samples, we can still recover the 21-cm signal, albeit with higher bias and uncertainty. This effect is predominantly observed in the region dominated by foregrounds ($k < 0.3 h cMpc^{-1}$), while the noise-dominated region remains relatively unaffected. Notably,




Figure 9. Ratios of recovered over input 21-cm power spectra using a VAE (top panel) and IAE (bottom panel) with different numbers of samples in the training set, ranging from 50 (blue) to 5000 samples (yellow). Each case is depicted with a small offset around the measured k -bin for clarity. The 2σ uncertainty is also indicated for each case.

the case with 500 samples performs nearly as well as that with 5000 samples.

These findings demonstrate the capability of our methods to achieve reliable recovery of the 21-cm signal, even with limited training data. While a larger training set generally leads to improved performance, our approach proves robust in cases where computational constraints limit the availability of extensive training samples.

5 DISCUSSION AND CONCLUSIONS

Detecting the redshifted 21-cm radiation of neutral hydrogen is the ultimate probe of the early stages of the observable Universe (Dark Ages and EoR). Even though a tremendous observational effort is ongoing – culminating with the SKA Observatory being online by the end of this decade – these observations are hard to perform as the signal is dwarfed by astrophysical contaminants shining at the same wavelengths, which are coupled to intricate systematics. Among the strategies developed to address these challenges, we find GPR, introduced in the 21-cm context by Mertens, Ghosh & Koopmans (2018) and that has become the method of choice of several 21-cm experiments.

The scope of this work has been improving GPR, making it more adaptable to the component separation task to retrieve the 21-cm signal. In particular, standard GPR assumes that the frequency–frequency covariance of the different components can be expressed analytically. We substituted that framework with covariance models learnt from 21-cm signal simulations. To do so, we relied on two different AE: variational and interpolatory AE. We tested the three methods (GPR, VAE–GPR, and IAE–GPR) in various scenarios to assess their ability to pinpoint the accurate power spectra of the 21-cm signal.

First, we demonstrated that using $n = 2$ latent variables is sufficient to characterize the spherically averaged power spectrum of the 21-cm signal using both the VAE and IAE schemes based on our study’s specific 21cmFast simulation set. Increasing the number of samples

in the training sets enhances the reconstruction capability of the AE, resulting in a representation that better reflects the true signal. However, it is important to note that this improvement in training does not necessarily translate to better inference when recovering the 21-cm signal from a simulated observation. Generally, the AE reconstruction is more successful for power spectra with flatter profiles compared to those with more complex features. In the best cases, both VAE and IAE perform similarly. Still, VAE is more prone to mismodel power spectra with intricate features.

We then demonstrated the advantages of employing ML–GPR for component separation tasks by evaluating various scenarios. The inherent k_{\perp} dependence in the learned kernel proved significant, distinguishing it from the standard GPR approach, which lacks this information. In general, both VAE and IAE-based kernels exhibited improved recovery of the 21-cm signal power spectra, characterized by reduced bias and lower uncertainty than standard GPR. This can be attributed to the learnt kernel being constrained to a set of covariances that align with our expectations for the 21-cm signal.

Promisingly, when an additional contaminant, simulating a systematic in the data, was introduced to the data cube, both VAE–GPR and IAE–GPR demonstrated their effectiveness. When the systematic was properly accounted for in the prior covariance model, these methods successfully recovered the power spectra of the 21-cm signal, although with higher uncertainty. Moreover, even when the systematic was not accounted for, VAE–GPR and IAE–GPR still provided a correct upper limit on the power spectra of the 21-cm signal, by not overfitting the unknown component. This highlights the robustness of VAE–GPR and IAE–GPR in the presence of systematics, showcasing their capability to deliver accurate signal estimates in challenging scenarios.

Significantly, the two ML–GPR outperform GPR in the scenario where we use a 21-cm signal not represented in the training set (nor in the GPR analytical kernel framework). The latter test is essential before applying ML–GPR in a real observational context, as it assesses the importance of the unknown representativeness of the training set.

In general and expectedly, uncertainties are higher in the k -regions where foregrounds or noise are dominant. When increasing the total observing time, we observed a decrease in uncertainty, as anticipated, in the noise-dominated region. However, the decrease in uncertainty was only marginal in the foreground-dominated region. We also found that for this specific 21cmFast simulation set, the number of samples in the training set does not play a significant role in the component separation task.

As for the differences between the two new algorithms, IAE–GPR tends to have a slightly higher fidelity recovery of the signal power spectra than VAE–GPR. We remind that VAE looks after a global description of the manifold where the 21-cm power spectra live, tending to generalize. However, IAE works with true anchor points on the manifold from which the algorithm can interpolate the output power spectra, making it more flexible in the reconstruction.

For real applications in the short term, we expect the different performances of the two ML–GPR to be insignificant as we deal with highly contaminated and noise-dominated data. None the less, the IAE architecture is robust also in the case of few available training samples, which, in the short term, can allow us to train it with 21-cm simulations that are more realistic than what we have exploited here, but also more computationally expensive and thus more scarce.

The 21-cm from the EoR and CD still needs to be better understood; many simulation codes are integrating more complex physics. We should combine multiple simulations code and train on

a combined set to widen the prior of the 21-cm component and cover all possibilities.

As another future outlook, we plan to further test these methods' robustness against more realistic contaminants. The test with the excess component (Scenario IV, left column of Fig. 7) shows that in the presence of an unaccounted component, its contribution to the variance in the data will be absorbed by other components – most likely by the 21-cm part which is flexible enough. As such, constructing the correct covariance prior model that can account for all data systematics is of utmost importance, informed in this stage by solid statistical techniques.

Moreover, we expect AE to be of use also for modelling the contaminant part of the observed signal, which is the foreground for which it is harder to obtain, for example, an analytical covariance, especially if we want to incorporate more realistic effects due to, for example, the beam and ionosphere. For that to be feasible, we first need to assemble more sensible simulations. The latter is part of our longer endeavor as we approach the SKAO era.

Summarizing, the critical outcomes of our work are that GPR with a trained kernel can recover a 21-cm signal that was not described in the training set and, to some extent, can do so also when an unknown component – which we will not be able to simulate – contaminates the data cube. Our results show the strength of these generative approaches and represent a further step forward in optimizing the analysis techniques that will make the detection at $z > 6$ of the 21-cm signal possible.

ACKNOWLEDGEMENTS

FGM acknowledges support of the PSL Fellowship. IPC is supported by the European Union within the Next Generation EU programme (PNRR-4-2-1.2 project No. SOE.0000136). The authors thank the anonymous referee for the constructive feedback that helps improve this manuscript.

DATA AVAILABILITY

The code to reproduce our results, as well as the 21cmFast training set and a step-by-step procedure, is available at https://gitlab.com/flomertens/ml_gpr_methodo_paper. Further code is also available at https://gitlab.com/flomertens/ps_eor.

REFERENCES

- Bobin J., Carloni Gertosio R., Bobin C., Thiam C., 2023, *Digit. Signal Process.*, 139, 104058
- Bonaldi A., Brown M. L., 2015, *MNRAS*, 447, 1973
- Chapman E. et al., 2012, *MNRAS*, 423, 2518
- Chapman E. et al., 2013, *MNRAS*, 429, 165
- Cheng C. et al., 2018, *ApJ*, 868, 26
- Ciardi B., Ferrara A., Marri S., Raimondo G., 2001, *MNRAS*, 324, 381
- Datta K. K., Choudhury T. R., Bharadwaj S., 2007, *MNRAS*, 378, 119
- Datta A., Bowman J. D., Carilli C. L., 2010, *ApJ*, 724, 526
- Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2017, *MNRAS*, 470, 1849
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publ. Astron. Soc. Pac.*, 125, 306
- Furlanetto S. R., 2016, in Mesinger A. ed., *Astrophysics and Space Science Library*, Vol. 423, *Understanding the Epoch of Cosmic Reionization: Challenges and Progress*. Springer International Publishing, Switzerland, p. 247
- Gehlot B. K. et al., 2019, *MNRAS*, 488, 4271
- Gehlot B. K. et al., 2020, *MNRAS*, 499, 4158
- Ghosh A. et al., 2020, *MNRAS*, 495, 2813

- Goodman J., Weare J., 2010, *Commun Appl. Math. Comput. Sci.*, 5, 65
- Hinton G. E., Salakhutdinov R. R., 2006, *Science*, 313, 504
- Hogg D. W., 1999, preprint (astro-ph/9905116)
- Jacobs D. C. et al., 2016, *ApJ*, 825, 114
- Kern N. S., Liu A., 2021, *MNRAS*, 501, 1463
- Kern N. S., Parsons A. R., Dillon J. S., Lanman A. E., Fagnoni N., de Lera Acedo E., 2019, *ApJ*, 884, 105
- Kingma D., Ba J., 2014, *Comput. Sci.*
- Kingma D. P., Welling M., 2013, preprint (arXiv:1312.6114)
- Kolopanis M. et al., 2019, *ApJ*, 883, 133
- Koopmans L. et al., 2015, *PoS, AASKA14*, 001
- Liu A., Shaw J. R., 2020, *Publ. Astron. Soc. Pac.*, 132, 062001
- 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
- Mertens F. G., Ghosh A., Koopmans L. V. E., 2018, *MNRAS*, 478, 3640
- Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
- Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Morales M. F., Bowman J. D., Hewitt J. N., 2006, *ApJ*, 648, 767
- Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, *ApJ*, 752, 137
- Murray S. G., Trott C. M., Jordan C. H., 2017, *ApJ*, 845, 7
- Offringa A. R., Mertens F., Koopmans L. V. E., 2019, *MNRAS*, 484, 2866
- Park J., Mesinger A., Greig B., Gillet N., 2019, *MNRAS*, 484, 933
- Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012, *ApJ*, 756, 165
- Patil A. H. et al., 2016, *MNRAS*, 463, 4317
- Patil A. H. et al., 2017, *ApJ*, 838, 65
- Pritchard J. R., Loeb A., 2012, *Rep. Prog. Phys.*, 75, 086901
- Rasmussen C. E., Williams C. K. I., 2005, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press
- Semelin B., Eames E., Bolgar F., Caillat M., 2017, *MNRAS*, 472, 4508
- Spinelli M., Carucci I. P., Cunnington S., Harper S. E., Irfan M. O., Fonseca J., Pourtsidou A., Wolz L., 2021, *MNRAS*
- Thyagarajan N. et al., 2015, *ApJ*, 804, 14
- Trott C. M., Wayth R. B., Tingay S. J., 2012, *ApJ*, 757, 101
- Trott C. M. et al., 2016, *ApJ*, 818, 139
- Trott C. M. et al., 2020, *MNRAS*, 493, 4711

APPENDIX A: IMPLEMENTATION DETAILS ABOUT THE VAE

Both the IAE and VAE used in this study are based on different implementations of MLP-based (standard multilayer perceptron) autoencoder architectures. Common parameters for both the architecture and the training are reported in Table A1. These architecture are conceptually different. The information bottleneck is implemented in a different way: In a VAE, the latent space is enforced to have a Gaussian distribution. In the case of the IAE, dimension reduction is

Table A1. Details of the architecture of both the IAE and VAE used in this work.

	VAE	IAE
No of layers	3	3
Encoder	[20, 20, 20]	[12, 12, 12]
Decoder	[20, 20, 20]	[12, 12, 12]
Batch size	128	64
Batch size	128	64
Epochs	2000	5000
Beta	1e-4	1e-3
Latent dim	2	3
Learning rate	0.02	0.01
Learning rate decay	10^{-5}	No learning rate decay
Minimiser	Adagrad	Adam
Implementation	Pytorch	JAX

explicitly ensured by projecting the encoded samples on to the affine hull of encoded so-called anchor points. In the present experiment, the number of anchor points is fixed to 3 and have been chosen at random from the training set. The loss functions of both AE are composed of two terms: a reconstruction fidelity term that favors reconstruction and a latent space regularization term. For the VAE, the regularization term enforces the distribution to be close to a Gaussian distribution. In the case of the IAE (see equation 14), it acts differently by enforcing the encoded training samples to be close to the affine hull of the encoded anchor points. In both cases, the parameter β controls the trade-off between these two terms. When β vanishes, the training loss boils down to the reconstruction term only.

APPENDIX B: TEST WITH A MORE ‘EXTREME’ MODEL

In this section, we conduct a test in which the input signal is deliberately chosen to fall outside the range covered by the power spectra in our training set. The purpose is to assess the generative capabilities of the VAE and IAE models. We simulated a 21-cm signal using 21cmFast with a virial temperature of $T_{\text{vir}} = 10^7$, which is beyond the usual range considered. We also conducted this test with a noise level equivalent to 1000 h of SKA, because this signal was of lower amplitude compared to the one considered until now. The results of this test for the VAE and IAE models are shown in Fig. B1. While the recovery is not perfect, the input signal falls within the 2σ uncertainty for almost all k -modes. This exercise gives a similar outcome as our test in Section 4.2.5.




Figure B1. Recovered 21-cm signal power spectra using different methods: VAE-GPR (blue) and IAE-GPR (orange) for a more ‘extreme’ model, with simulation parameters outside the range typically considered. The power spectra of the noise component are shown in yellow, and the dashed black line represents the power spectra of the 21-cm component. The shaded area represents the 2σ uncertainty on the recovered 21-cm signal. The top panel displays the spherically averaged power spectra, while the center and bottom panels show the cylindrically averaged power spectra as a function of k_{\parallel} and k_{\perp} (averaged over k_{\perp} and k_{\parallel} , respectively).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.