



## Machine Learning and the SKA for Cosmic Dawn and the Epoch of Reionization

Anshuman Acharya,<sup>7,13</sup> Michele Bianco,<sup>9,10</sup> Daniela Breitman,<sup>11</sup> Huaxi Chen,<sup>18</sup> Abhirup Datta,<sup>1</sup> Kangning Diao,<sup>13,14</sup> Sambit K. Giri,<sup>2,3</sup> Caroline S. Heneka,<sup>12</sup> Nicholas Kern,<sup>4</sup> Adrian Liu,<sup>19</sup> Yashrajsinh Mahida,<sup>1</sup> Suman Majumdar,<sup>1</sup> Samit Kumar Pal,<sup>1</sup> Shulei Ni,<sup>18</sup> Yannic Pietschke,<sup>12</sup> Davide Piras,<sup>5</sup> Abinash Kumar Shaw,<sup>6,7</sup> Hayato Shimabukuro,<sup>15,16,17</sup> Ce Sui,<sup>14</sup> Anshuman Tripathi<sup>1</sup> and Xiaosheng Zhao<sup>8</sup>

<sup>1</sup>Department of Astronomy, Astrophysics & Space Engineering, Indian Institute of Technology Indore, Indore 453552, India

<sup>2</sup>Department of Astronomy and Oskar Klein Centre, AlbaNova, Stockholm University, SE-10691 Stockholm, Sweden

<sup>3</sup>Van Swinderen Institute for Particle Physics and Gravity, University of Groningen, Nijenborgh 3, 9747 AG Groningen, The Netherlands

<sup>4</sup>Department of Physics, University of Michigan, 450 Church St., Ann Arbor, MI, 48109-1040, USA

<sup>5</sup>Département de Physique Théorique, Université de Genève, 24 quai Ernest Ansermet, 1211 Genève 4, Switzerland

<sup>6</sup>Department of Computer Science, University of Nevada Las Vegas, 4505 S. Maryland Pkwy., Las Vegas, NV 89154, USA

<sup>7</sup>Max-Planck-Institut für Astrophysik, Garching D-85748, Germany

<sup>8</sup>Department of Physics & Astronomy, The Johns Hopkins University, Baltimore, MD 21218, USA

<sup>9</sup>Laboratoire d'Astrophysique, École Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, Chemin Pegasi 51, CH-1290 Versoix, Switzerland

<sup>10</sup>Institute for Particle Physics and Astrophysics, ETH Zurich, Wolfgang-Pauli-Str 27, CH-8093 Zurich, Switzerland

<sup>11</sup>Scuola Normale Superiore (SNS), Piazza dei Cavalieri 7, Pisa, PI, 56125, Italy

<sup>12</sup>Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany

<sup>13</sup>Berkeley Center for Cosmological Physics, University of California, Berkeley, CA 94720, United States

<sup>14</sup>Department of Astronomy, Tsinghua University, Beijing 100084, China

<sup>15</sup>South-Western Institute for Astronomy Research (SWIFAR), Yunnan University, Kunming, Yunnan 650500, People's Republic of China

<sup>16</sup>Key Laboratory of Survey Science of Yunnan Province, Yunnan University, Kunming, Yunnan 650500, People's Republic of China

<sup>17</sup>Nagoya University, Graduate School of Science, Division of Particle and Astrophysical Science, Chikusa-Ku, Nagoya, 464-8602, Japan

<sup>18</sup>Research Center for Computational Earth and Space Science, Zhejiang Laboratory, Hangzhou 311121, China

<sup>19</sup>Department of Physics and Trottier Space Institute, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

E-mail: [adrian.liu2@mcgill.ca](mailto:adrian.liu2@mcgill.ca), [mbianc@phys.ethz.ch](mailto:mbianc@phys.ethz.ch),  
[heneka@thphys.uni-heidelberg.de](mailto:heneka@thphys.uni-heidelberg.de)

When operational, the SKA will generate unprecedented amounts of data and provide exquisite sensitivity for 21 cm tomography of Cosmic Dawn (CD) and the Epoch of Reionization (EoR). With this comes opportunities for new data-driven algorithms that unlock new methods for instrument modelling, data analysis, theoretical simulation, and inference for understanding the high-redshift universe. In this chapter, we provide an overview of some machine learning algorithms that have been proposed for CD and EoR science with the SKA

## 1 Introduction

The Square Kilometre Array (SKA) promises to revolutionize the study of Cosmic Dawn (CD) and the Epoch of Reionization (EoR). A series of impressive precursor experiments (such as PAPER, MWA, LOFAR, HERA) have set an increasingly stringent set of upper limits on 21 cm emission from these epochs, and the SKA provides the necessary sensitivity to not only *detect* but also to *characterize* the cosmological signal to high significance. By simply using conventional data analysis techniques to produce conventional summary statistics such as the power spectrum, the SKA's unprecedented sensitivity is already expected to deliver world-leading constraints on the astrophysics governing CD and EoR.

Going beyond quantitative differences, the SKA pushes 21 cm studies of CD and EoR into qualitatively new regimes. SKA data rates, for example, will be larger than for any radio telescope to date and will necessitate real-time processing that has largely been an optional convenience for current observatories. In terms of science products, the SKA is expected to move beyond summary statistics and to produce images of CD and EoR, enabling a new and incisive set of studies.

With new challenges come new opportunities, and in this chapter we examine the way in which the development of the SKA has been coincident with another fast-developing field—machine learning (ML). Given the relatively recent rise of artificial intelligence techniques in astronomy, it would be fair to say that it is as-yet unclear what the ultimate impact of ML on the SKA will be. However, in this chapter we provide a speculative look at a sampling of ideas that have already circulated in the literature. These ideas span the full breadth of activities of the SKA, from low-level data analysis algorithms to new ways to extract CD and EoR science.

## 2 Observations

A key obstacle to the detection and characterization of CD and EoR is in understanding the interplay between instruments and sky emission. In other words, what is crucial is the understanding of the cosmological *as seen by the instrument*. This has led to efforts both in simulating instruments to high precision and in the development of novel analysis methods. In this section, we examine the various areas in which machine learning has led to new ways of bridging the instrument-sky divide.

### 2.1 Radio Frequency Interference

Radio Frequency Interference (RFI) represents a major obstacle to the detection and characterization of the cosmological 21 cm signal and other faint radio sources targeted by the SKA. RFI arises from a wide variety of terrestrial and satellite-based transmitters, and its highly variable nature across both time and frequency domains can mimic or obscure genuine astronomical signals. Effective mitigation of RFI contamination is therefore essential for preserving the scientific integrity of the data and enabling precision measurements of cosmological and astrophysical phenomena.

A number of traditional approaches have been developed to identify and suppress RFI in radio observations. Early techniques relied on the statistical detection of outliers in the time–frequency domain, employing algorithms such as SumThreshold (Offringa et al., 2010) and CUMSUM (Baan et al., 2004). These methods typically detect and flag RFI by identifying significant deviations

from local noise statistics and have been implemented in widely used software such as AOFlagger (Offringa, 2010). Other approaches have exploited statistical estimators such as spectral kurtosis (e.g., Nita and Gary, 2010) to detect non-Gaussian signal components, or applied baseline fitting and harmonic analysis to suppress large-scale or periodic interference (e.g., Maan et al., 2021). While these classical techniques are reliable for specific RFI morphologies, they are often constrained by heuristic thresholds and limited sensitivity to complex or evolving interference patterns.

With the rapid growth of data-driven astronomy, deep learning (DL) models have emerged as powerful alternatives that can operate directly on preprocessed spectrograms or visibility data, reducing the need for manual feature engineering. The U-Net architecture (Akeret et al., 2017), originally designed for biomedical image segmentation, was among the first deep models applied to RFI detection. Building on this foundation, subsequent studies introduced variants such as RFI-Net (Yang et al., 2020), which incorporated residual connections, and DSC-Dual-ResUNet (Zhang et al., 2021), which employed dual U-Net structures with depthwise separable convolutions to enhance efficiency. Pagano et al. (2023) used a U-Net architecture to tackle an RFI-related task: after RFI flagging, some analysis treatments *in-paint* the resulting gaps with statistically representative synthetic data. Pagano et al. (2023) examined the robustness of using U-Nets to performing inpainting, performing a comparison with other techniques.

Other works have explored recurrent neural networks (RNNs) and long short-term memory (LSTM) layers to capture temporal correlations in dynamic RFI signals (Burd et al., 2018). Beyond architectural innovations, Vafaei Sadr et al. (2020) also investigated alternative training strategies, including transfer learning, to enhance model performance and generalization. Collectively, these efforts demonstrate the ability of DL architectures to learn complex RFI patterns directly from large datasets, achieving superior accuracy and adaptability compared to traditional methods.

Beyond supervised learning, several groups have pursued unsupervised and generative approaches to address the scarcity of labeled RFI data. Generative adversarial networks (GANs) have been trained to model uncontaminated background data and identify anomalies corresponding to RFI events (Du Toit et al., 2024), while convolutional autoencoders (CAEs) have been used to learn latent representations of clean data for anomaly detection (Mosiane et al., 2016). In one example, Mesarcik et al. (2022) combined CAE-based latent representations with KNN clustering to identify RFI in contaminated observations. Similarly, clustering algorithms such as DBSCAN have been applied to group false-positive events in reduced feature spaces, effectively identifying outlier patterns associated with interference (Jacobson-Bell et al., 2025). These approaches highlight the growing shift toward data-driven, model-agnostic strategies capable of adapting to diverse observational conditions without explicit supervision.

Despite these advances, several challenges remain. ML-based models—particularly deep architectures—can be sensitive to the domain and distribution of their training data. Their performance often degrades when applied to datasets with differing instrumental configurations, observing conditions, or spectral resolutions (Yang et al., 2020; Du Toit et al., 2024). Furthermore, the deployment of ML methods in large-scale pipelines raises practical concerns about computational efficiency, interpretability, and the need for continual retraining as the RFI environment evolves.

While ML methods are promising, most operational pipelines still rely primarily on classical statistical flaggers, with ML approaches largely remaining at the validation or complementary stage. Continued progress will depend on developing domain-adaptive, interpretable, and computationally efficient ML architectures that can be seamlessly integrated into the SKA’s real-time data processing framework.

## 2.2 Ionospheric Effects

The ionosphere, a turbulent, low-density, multispecies ion plasma that surrounds the Earth, is driven mainly by solar extreme ultraviolet radiation and is a known source of direction-dependent effects (DDEs) in low-frequency radio astronomy. These DDEs introduce systematic errors that fundamentally limit on the dynamic range of radio interferometric observations, particularly for experiments targeting the CD-EoR signal. For a comprehensive discussion of the mitigation of ionospheric errors with SKA-Low, we refer the reader to the [SKA foregrounds chapter \(2025\)](#). The dynamic nature of the ionosphere produces flicker noise ([Datta et al., 2014](#)) ( $1/f$  noise, where  $f$  is the dynamical frequency). These errors do not reduce with long integration.<sup>1</sup> Hence, there is a need for ionospheric calibration at shorter time intervals that match the ionospheric variability. The residual DDEs or the residual ionospheric refractive shifts have, in some studies, been a non-negligible component in residual excesses. the upper limits of the CD/EoR power spectrum measurements ([Pal et al., 2025](#)). Therefore, mitigating these residual DDEs is essential for robust interpretation of the faint CD-EoR signal in both the Fourier and image domains. For example, a recent study by [Brackenhoff et al. \(2024\)](#) demonstrated that the excess noise induced by ionospheric phase errors can be removed using Gaussian Process Regression (GPR; see [Sec. 2.5](#) for more details and another application). They found that the propagation of ionospheric errors from longer to shorter baselines, which are the most affected by induced phase errors, may not be a dominant contributor to the excess noise in this LOFAR study.

In [Tripathi et al. \(2024\)](#), an artificial neural network (ANN)–based framework was developed to recover the global 21-cm signal in the presence of foreground and ionospheric effects. The study found that, for slowly varying ionospheric conditions, a simple ANN architecture is capable of successfully recovering the underlying cosmological signal.<sup>2</sup> Several machine learning frameworks have also been developed to reconstruct the Total Electron Content (TEC) in radio interferometry calibration, correcting the DDEs arising from ionospheric effects ([Li et al., 2025](#); [de Jong et al., 2025](#)). [Albert et al. \(2020\)](#) propose a probabilistic physics-informed model using Gaussian processes to infer ionospheric phase screens, thereby reducing residual errors.

## 2.3 Gain Calibration

The gain calibration errors are expected to be present in any radio interferometric observations of the CD-EoR. Generally, the calibration is done by estimating the Jones matrices with the assumption

---

<sup>1</sup>Some reduction will naturally result from stacking multi-night data, which can cause a decorrelation between night-to-night independent components. However, systematic effects (e.g., calibration errors) may still be present.

<sup>2</sup>In essence, the ANN is taking the place of the usual Bayesian inference exercise of fitting a global signal or a power spectrum to infer astrophysical parameters. While currently it is still generally practical to perform Bayesian inference once a summary statistic has been produced, future analyses that incorporate yet more nuisance parameters may benefit from the faster runtime of a machine learning-based algorithm.

of a perfectly known model sky, or (in cases where a radio interferometer has many repeated copies of the same baselines) by exploiting the self-consistency of a redundant array. However, limitations such as parameter degeneracies, incompleteness of sky models, imperfect knowledge about the primary beam, and rapid atmospheric fluctuations result in residual calibration errors. These residual errors can lead to incorrect interpretations of the observed signal (Datta et al., 2009; Beardsley et al., 2016), introducing spurious features in the image domain (Pal et al., 2025) and biasing statistical estimates of the signal in the Fourier domain, such as the power spectrum (Barry et al., 2016; Kern et al., 2020; Byrne et al., 2021; Mazumder et al., 2022). Such biases can propagate through the inference pipeline to the estimation of astrophysical parameters, and thus ultimately lead to misinterpretation of the underlying physical processes (Tripathi et al., 2025). For estimators of signal power spectrum, these calibration errors can introduce excess power across different  $k$ -modes.

Mitigating these errors is therefore crucial for a robust interpretation of the signal in both image and Fourier domains. Motivated by this, various methodologies have been developed for this purpose. For instance, to address the excess power introduced in the power spectrum by residual gain errors, techniques such as Gaussian Process Regression (GPR) and Principal Component Analysis (PCA) along with other unsupervised machine learning methods have been applied to identify the  $k$ -modes within the EoR window most affected by calibration systematics, enabling the construction of more reliable estimators for the signal power spectrum (Mertens et al., 2020; Chen et al., 2023; Brackenhoff et al., 2024; Beohar et al., 2025). A recent study by Beohar et al. (2025) demonstrates that a hybrid strategy combining avoidance with PCA and GPR enables robust signal recovery across the entire  $k$ -range. They found that GPR and PCA-based gain error correction performs more effectively at large scales, although both methods tend to underestimate the power on smaller scales.

In addition, deep learning-based models, particularly U-Net architectures, have been employed to correct for gain errors arising from frequency-dependent Gaussian beams and bandpass fluctuations, thereby improving 21-cm signal recovery (Chen et al., 2024). Leeney et al. (2025) proposed a machine learning-based calibration framework for a 21-cm sky-averaged cosmology experiment, aimed to accurately model the complex gain behaviour of the radio instrument.

## 2.4 Imaging

One of the significant innovations of the SKA-Low telescope is that it will be sensitive enough not only to detect the 21-cm signal but also to produce 3D tomographic maps of its distribution at different stages of the EoR (Wyithe et al., 2015; Mellema et al., 2015; Giri et al., 2018b). For a comprehensive discussion of 21-cm signal imaging with SKA-Low, we refer the reader to SKA imaging chapter (2025). Over the years, several AI methods have been developed to process images and 3D data (e.g. Le Cun et al., 1997; Baldi, 2012; Ronneberger et al., 2015; Kingma and Welling, 2022), so naturally, machine learning techniques have been establishing their role as the primary method for image analysis among cosmologists and astrophysicists (see Dvorkin et al., 2022; Lahav, 2023, for a review). The application of these techniques to 21-cm cosmology has accelerated in recent years, driven by the need to extract maximum scientific information from the complex, high-dimensional tomographic datasets that SKA-Low will produce. In this section, we focus on

the imaging process itself, reserving *applications* of imaging to later sections.

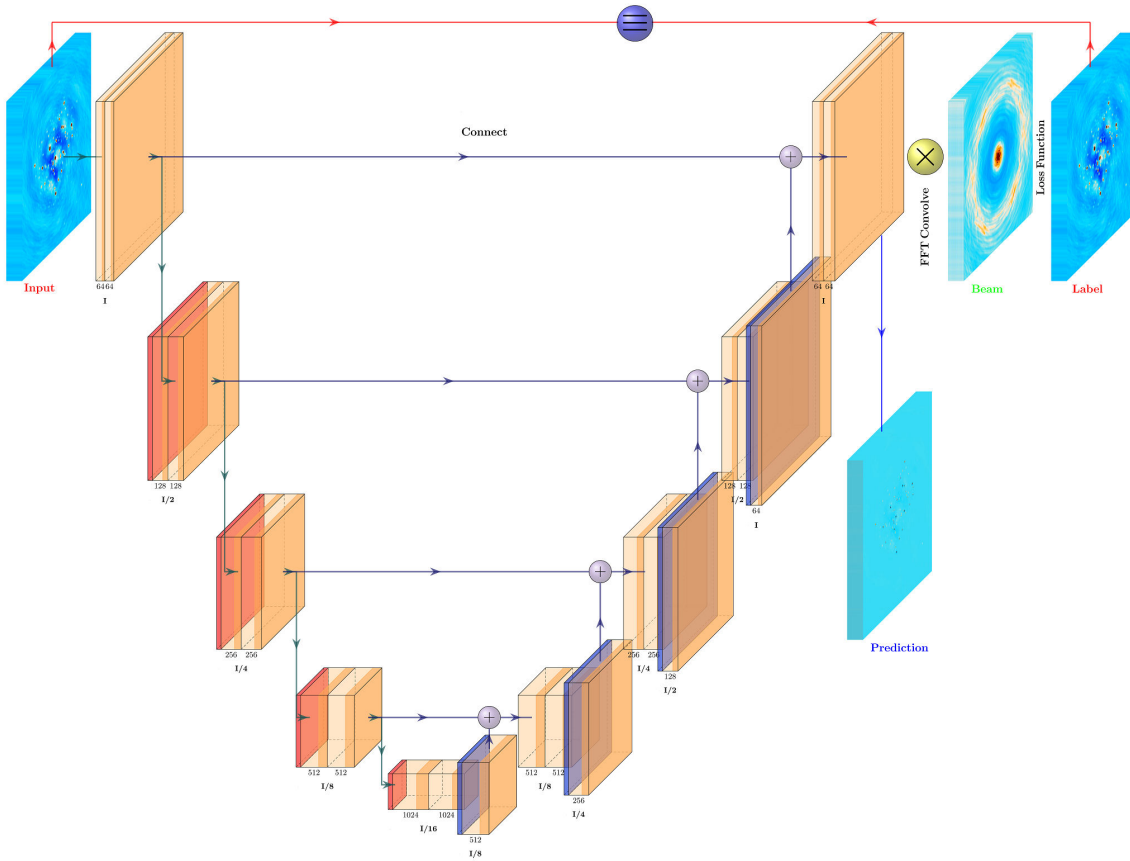
A fundamental challenge is addressing the systematic errors in radio interferometric imaging that arise from incomplete  $uv$ -coverage and, relatedly, the effect of the synthesised beam. These introduce artefacts in radio images that can be helpful to remove (“deconvolution”) for many applications, such as mitigating foreground contamination from extragalactic compact or point-like radio sources. Traditionally, CLEAN-based methods have been used for radio interferometric imaging. Initially proposed by Högbom (1974), CLEAN iteratively deconvolves the telescope beam by identifying and subtracting point sources from the dirty map, with the accumulated model convolved with a restoring beam to form the final image. Over the years, extensions of this method have been suggested, such as the multi-scale, multi-frequency model, implemented in MS-CLEAN and MF-CLEAN, for wide-band image convolution (Rau and Cornwell, 2011), and wide-field techniques such as  $w$ -projection and faceting, implemented in WSClean (Offringa et al., 2014), correct for direction-dependent effects. Although effective for compact sources, these methods struggle with diffuse or complex foregrounds, as representing sky brightness with discrete components can introduce additional artefacts. Their iterative nature also makes them computationally demanding, especially for wide-field, low-frequency surveys. An unsuccessful deconvolution process (or at least one whose statistical properties are not well-understood) may exacerbate foreground contamination and increase the difficulty of signal separation (Hothi et al., 2020; Ni et al., 2022; Wang et al., 2025; Cox et al., 2024). This, therefore, motivates the exploration of alternative deconvolution approaches.

To overcome some of the limitations mentioned above, researchers have proposed various AI-based deconvolution approaches. For instance, Schmidt et al. (2022) and Geyer et al. (2023) employ a convolutional neural network to reconstruct images of compact radio sources by directly inferring missing visibility data from incomplete Fourier-space data, offering improvements in speed, automation, and reproducibility over traditional methods such as CLEAN. Another noteworthy example is PI-AstroDeconv (Ni et al., 2024, 2025), whose network architecture is illustrated in Figure 1. Rather than relying solely on deep networks for blind learning, PI-AstroDeconv constructs an explicit imaging model based on the physical principles of observational imaging:

$$I_D(x, y) = P_D(x, y) \otimes I'(x, y), \quad (1)$$

where the dirty map  $I_D(x, y)$  is the convolution of the clean map  $I'(x, y)$  with the telescope beam function  $P_D(x, y)$ , thereby explicitly modelling the effect of the beam on the observed image, as illustrated in Figure (1). Based on this model, PI-AstroDeconv transforms the beam deconvolution problem into an image-to-image mapping task: given a dirty map and its corresponding PSF, the network directly outputs an estimated clean map, which is then convolved with the same PSF to produce a “dirty map” for loss computation. The training minimises only the Log-Cosh distance between simulated and true dirty maps, achieving fully unsupervised learning without requiring clean images as labels.

PI-AstroDeconv operates directly on image-domain data; its inputs are the dirty maps produced by interferometric imaging and the corresponding PSFs, without relying on conventional deconvolution operations. This design substantially simplifies the preprocessing workflow and introduces minimal



**Figure 1:** The PI-AstroDeconv architecture. The yellow blocks represent convolutional layers, red blocks represent pooling layers, and gray-green blocks represent upsampling operations. The left half of the network depicts the downsampling path, while the right half represents the upsampling path. The number of channels is indicated at the bottom of each block. The arrows above the image symbolize the connections between network layers. The numbers below each layer in the network indicate the quantity of convolutional kernels present in that specific layer. Additionally, the letters denote the variations in the image dimensions within the layer. The letter "I" specifically represents the size of the input image. The blocks located at the far left and far right of the network indicate the input and the label, respectively (represented as deep blue circles). A Depthwise convolutional layer, represented as a yellow circle, is inserted between the U-Net network and the label. This layer utilizes convolutional kernels resembling telescope beams and is accelerated using FFT. The U-Net's output is convolved with the beam, and the result is compared with the labels (Ground Truth) to compute the loss function. Network inputs and ground truth labels are sourced from the telescope. The model is trained to embed physical priors through convolutional layers, ensuring outputs align with the original observations. The final prediction of the network is derived from the last layer of the U-Net, represented as the last yellow layer. The symbols +,  $\times$ , and  $\equiv$  represent concatenation, convolution, and exact equality, respectively.

additional computational overhead. To accommodate the large size and high resolution of radio images, the network performs the final convolution in the frequency domain using FFTs, reducing computational complexity from  $\mathcal{O}(n^4)$  to  $\mathcal{O}(n^2 \log n)$ , significantly lowering memory usage and training time. The convolution kernel is fixed as the input PSF and does not update with gradients, embedding the telescope beam response as a hard constraint in the loss function to ensure physical

consistency.

## 2.5 Foreground Mitigation

A key challenge in 21-cm imaging is that astrophysical foregrounds (e.g., synchrotron and free-free emissions from the Milky Way and other galaxies) dominate over the cosmological signal in a wedge-shaped region of the cylindrical Fourier space ( $k_{\perp}$ ,  $k_{\parallel}$ ). One way to mitigate foreground is to filter out the wedge region (“foreground avoidance”) from the visibility data before image deconvolution. With this technique, we can focus on areas in the Fourier space that are mostly foreground-free (see Kerrigan et al., 2018; Morales et al., 2019; HERA Collaboration et al., 2023, for more details). However, this process removes a relevant portion of cosmological information along with the undesired foregrounds. Thus, a worthwhile goal is to reconstruct the 21-cm signal in the wedged region and/or successfully model and remove foregrounds, thereby maximising the physical scales that can be explored with the 21-cm power spectrum and boosting sensitivity by orders of magnitude (Pober et al., 2014).

Machine learning methods have been developed to recover the spatial information lost during wedge filtering by exploiting the non-Gaussian correlations in the 21-cm signal (Gagnon-Hartman et al., 2021; Kennedy et al., 2024). Deep generative models based on stochastic interpolants have also been introduced to reconstruct 21-cm data lost to wedge filtering, leveraging the non-Gaussian nature of the signal to map wedge-filtered 3D lightcones to samples from the conditional distribution of wedge-recovered lightcones (e.g. Sabti et al., 2024; Chen et al., 2025). Additionally, approaches have been developed to combine standard foreground mitigation strategies with novel AI techniques to recover complete 21-cm signal images directly from foreground-contaminated data. For instance, the work of Bianco et al. (2024) combines U-Nets with existing foreground removal methods (e.g., PCA, wedge filtering, polynomial decomposition, and GPR), improving the recovery of both neutral (HI) and ionised (HII) regions from tomographic data. Furthermore, the recovery of the 21-cm signal in tomographic maps has been enhanced by modifying the U-Net architecture to accept an additional input, which serves as prior information. With the addition of a *convolutional intercepting layer*—a series of pooling and convolutional operations that intercept and combine with the output of the skip connections that process images with residual foreground—additional maps are processed to integrate prior information on the neutral/ionised in observed foreground contaminated maps (Bianco et al., 2025), demonstrating an enhancement in the recovery of 21-cm tomographic data at small scales and across a wide range of astrophysical parameters and redshifts. These maps can be derived from observations of high-redshift galaxies.

Another approach to machine learning-inspired foreground mitigation has been to use Gaussian Process Regression (GPR). GPR can be employed to model and subtract foreground contamination (Mertens et al., 2018; Gehlot et al., 2019; Hothi et al., 2021), with each data component being modelled using analytic covariance kernels. However, the use of simplified analytic kernels for the EoR signal can limit the accuracy of component separation. It may bias the recovered 21-cm signal (for example, see Kern and Liu, 2021), motivating the development of more flexible, data-driven physics-motivated kernel models.

Machine learning offers a powerful alternative by learning covariance priors directly from physics-

based simulations. In particular, ML-enhanced GPR (ML-GPR) has been recently developed to enhance foreground mitigation (Mertens et al., 2024; Acharya et al., 2024a). This approach utilises unsupervised generative models (e.g., Variational Autoencoders) trained on simulated 21-cm lightcones from codes such as 21cmFAST (Mesinger and Furlanetto, 2007; Mesinger et al., 2011) or GRIZZLY (Ghara et al., 2015; Ghara et al., 2018) to build flexible, data-driven covariance kernels. These learned quantities then replace the fixed analytic kernels in the usual GPR pipeline. By fitting the latent parameters of the ML-derived kernel to the data, the method mitigates mismatches between assumed and true signal covariance. It can adapt to complex spectral features and non-Gaussianity of the 21-cm signal that rigid kernels miss. According to Acharya et al. (2024a), ML-GPR outperforms standard GPR in component separation and remains robust even when the true signal lies outside the training set, thereby providing resilience to instrumental systematics and noise. This has now been incorporated in foreground mitigation pipelines for various observations, such as with LOFAR (applied to the NCP field in Acharya et al. 2024b; Mertens et al. 2025, the 3C196 field in Ceccotti et al. 2025) and NenuFAR (Munshi et al., 2024, 2025). Additionally, it has demonstrated the potential to detect the 21-cm power spectrum by isolating the cosmological component not only from foregrounds but also from more complex systematics. In this regard, upper limits on the 21-cm power spectrum, after separating other contributions, have been reported in Acharya et al. (2024b) and Mertens et al. (2025).

A key innovation in ML-GPR is the use of a Variational Autoencoder (VAE; Kingma and Welling, 2013) to learn a latent representation of the 21-cm signal’s covariance structure. A VAE is a generative probabilistic model where observed data  $x$  (here, realizations of the 21-cm signal from simulations) are compressed into latent variables  $z$  drawn from a prior distribution  $p(z)$ , typically constrained to be as close as possible to a multivariate Gaussian  $\mathcal{N}(0, I)$ . The model consists of two neural-network components: an *encoder*,  $q_\phi(z|x)$  with parameters  $\phi$ , which approximates the intractable posterior distribution over the latent variables, and a *decoder*,  $p_\theta(x|z)$  with parameters  $\theta$ , which reconstructs data from latent samples. Training is achieved by maximizing the evidence lower bound (ELBO):

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)] - D_{\text{KL}}(q_\phi(z|x) \parallel p(z)). \quad (2)$$

In Eq. 2, the first term is the expected log-likelihood of reconstructing the data under the generative model, encouraging accurate recovery of  $x$  from its latent representation. The second term is the Kullback–Leibler divergence, which penalizes deviations of the approximate posterior  $q_\phi(z|x)$  from the prior  $p(z)$ . This regularization enforces smoothness and continuity in the latent space, preventing overfitting and enabling meaningful interpolation between latent variables.

In the ML-GPR framework (Mertens et al., 2024; Acharya et al., 2024a), only the encoder is retained after training. The encoder provides a mapping from high-dimensional 21-cm signal realizations into a structured latent space that captures their statistical correlations. From this representation, one can construct a data-driven covariance kernel,

$$k_{\text{ML}}(x, x') \approx \mathbb{E}_{z \sim q_\phi(z|x, x')} \left[ (x - \mu_z)(x' - \mu_z)^\top \right], \quad (3)$$

where  $\mu_z$  denotes the mean of the encoded latent distribution. This learned kernel  $k_{\text{ML}}$  replaces hand-crafted analytic kernels in GPR. In practice, the VAE encoder learns the manifold of possible

EoR signals from simulations, such that the covariance structure naturally reflects both astrophysical variability (e.g., ionization histories, source models) and instrumental effects (e.g., chromaticity, mode-mixing). This enables ML-GPR to remain flexible to realistic deviations, in contrast to standard GPR approaches where the covariance must be specified a priori.

In summary, ML-GPR leverages modern generative models to extend classical GPR for 21-cm analysis. By learning covariance kernels from simulations reduces prior mismatch and signal loss, adapts to non-ideal spectral features, and yields more accurate 21-cm reconstructions. These advances have already improved foreground pipelines in deep EoR fields, and further improvements (learnt kernels for the foregrounds, other ML algorithms, etc. apart from improvements in data reduction) are also in development. All these improvements promise a robust path toward detecting the cosmic 21-cm signal with the SKA-Low. Indeed, in the recent SKA Science Data Challenge 3a, foreground (SDC3a<sup>3</sup>), participants were asked to recover the 21-cm cylindrical power spectra from foreground contamination and systematics from visibility and/or tomographic data, and GPR resulted in the smallest residuals and least bias (Bonaldi et al., 2025).

### 3 Theory

In addition to experimental concerns, machine learning methods have also had an impact on theoretical calculations. In many ways, CD/EoR represent interesting regimes to apply machine learning, since (unlike many other sub-fields in cosmology) there are both computational limitations and considerable uncertainties as to the correctness of models.

#### 3.1 Accelerating & Improving Simulations at the Field Level

Numerical simulations are a key tool for studying the Epoch of Reionisation (EoR). Yet, reionisation is an intrinsically multi-scale problem: simulations must encompass cosmological volumes of order  $\sim 500$  Mpc to capture the large-scale topology of ionised regions (Iliev et al., 2014), while simultaneously resolving sub-galactic processes that shape the ionising sources (Trebtsch et al., 2017), over an extended period of cosmic history from  $z \sim 20$  to 5 (Giri et al., 2024). Meeting these requirements makes such simulations computationally expensive, demanding billions of CPU and/or GPU hours on thousands of nodes in the world’s most powerful supercomputers (e.g. Ocvirk et al., 2016; Ocvirk et al., 2020; Trebtsch et al., 2021; Garaldi et al., 2022). To mitigate these limitations, several EoR simulations have been adapted to run on GPU devices, enabling significant performance improvements and speed-ups (Aubert and Teyssier, 2010; Aubert et al., 2015; Hirling et al., 2024). However, despite these advances, such simulations remain computationally demanding. For this reason, recent studies have begun exploring AI techniques to overcome the computational constraints inherent in EoR simulations.

When integrating AI with EoR simulations, two main objectives can be identified. The first focuses on replacing the radiative transfer solver with a learned model that emulates its output, predicting quantities such as the hydrogen neutral fraction or ionisation field from input data describing the ionising sources (e.g. positions, intrinsic luminosities, and UV mean-free paths), the IGM density

---

<sup>3</sup><https://sdc3.skao.int/challenges/foregrounds>

distribution, or directly from a set of astrophysical parameters. The second aims to accelerate specific components of the simulation pipeline by employing neural networks that learn the complex relationships between astrophysical processes (e.g. star formation history, stellar radiative feedback) and their host environments (i.e. dark matter haloes and/or overdensities in the IGM) on scales that are generally not resolved in reionisation simulations.

An example of the first application was presented by [Chardin et al. \(2019\)](#), who employed a convolutional autoencoder that takes 2D images of the gas density field and the stellar number density from a simulated lightcone and reconstructs the reionisation time,  $t_{\text{reion}}(\mathbf{r})$ —the cosmological time at which a region reaches a given ionisation fraction threshold. Two separate encoders process the two inputs, and an averaging operation combines the two low-dimensional latent spaces before deconvolutional operations are applied. This method shows promising results for large-scale features ( $\geq 10$ , cMpc) but struggles to recover small-scale structures and tends to delay the appearance of the first ionised regions at high redshift. In contrast, [Korber et al. \(2023\)](#) proposed a different approach that focuses on recovering small-scale details of reionisation. In this work, a CNN processes input data from sub-volumes of a larger EoR simulation and predicts the neutral fraction of the central voxel. The training is enforced voxel-by-voxel using a physically informed loss function that approximates the hydrogen chemistry equation. The method employs the same assumptions and input fields, but introduces an additional preprocessing step that generates a smoothed source field based on an empirical redshift-dependent model of the mean free path of ionising photons,  $\lambda_{\text{mfp}}(z) \propto \frac{c}{H(z)}(1+z)^{-2.55}$ . The resulting field is then used as an extra input to the network. This approach achieves high accuracy on small scales but shows reduced large-scale performance at high redshift due to the mean-free-path assumption. Neither of the methods described above imposes a specific EoR model during training. Their central assumption is that a neural network can learn the underlying physical relationship between the sources and sinks of radiation and the resulting ionised field. In this sense, the network maps the redshift evolution of the ionisation front from the given source distribution, providing an implicit, data-driven approximation of radiative transfer based on pre-run numerical simulations.

Following a similar motivation, [Posture et al. \(2025\)](#) has integrated U-Net with a Vision Transformer-based architecture (CosmoUiT) to emulate the 3D ionisation field and the corresponding 21-cm field from the EoR. It takes 3D dark matter and halo fields, along with reionisation model parameters, as inputs. In this architecture, the multi-head self-attention mechanism of the transformer captures long-range dependencies, while the convolutional layers in the U-Net capture small-scale variations in the field. Furthermore, the reionisation parameters are provided during training to both the vision transformer block and the U-Net block, conditioning the output on these parameters. The hybrid design of the emulator overcomes the limitations of earlier approaches, capturing both large- and small-scale features of the 21-cm signal, and limiting the error at the boundary between ionised and neutral regions that leads to over-estimation of bubble size distribution at the early stage of reionisation.

Another approach consists of generating 21-cm tomographic maps directly from a set of astrophysical and cosmological parameters. These methods typically assume a specific model that parametrises the sources of radiation and recombination in the IGM, and are generally trained on

large datasets ( $\sim 10^4$  realisations) to cover the broad parameter space. As such, they can be regarded as emulators trained on a particular model of reionisation and are often integrated into pipelines for EoR parameter inference. Examples include [Chardin et al. \(2019\)](#); [Korber et al. \(2023\)](#); [Zhao et al. \(2023\)](#); [Hothi et al. \(2025\)](#); [Posture et al. \(2025\)](#); [Heneka et al. \(2025\)](#). [Zhao et al. \(2023\)](#) employs two kinds of generative models to emulate 21-cm image slices: Generative Adversarial Networks (GANs, [Goodfellow et al., 2014](#)) and diffusion models ([Sohl-Dickstein et al., 2015](#); [Ho et al., 2020](#)). While GANs are more computationally efficient, diffusion models have been shown to produce a greater diversity of samples, as quantified by scattering transform coefficients. The main limitation of this approach is that 21-cm slices in the training set are limited to high redshifts, where the 21-cm signal exhibits little evolution and remains close to a Gaussian density field. As a result, the generative models often struggle to learn meaningful parameter dependencies from such data. The work of [Heneka et al. \(2025\)](#) employs large language models (LLMs) to analyse and generate 21-cm tomographic maps from astrophysical and cosmological parameters. In this study, pre-trained transformers (e.g. Qwen2.5) are adapted for out-of-domain physical inference tasks, highlighting an intriguing analogy between natural language and cosmological structure formation. The model accepts a sequence of  $\delta T_b$  maps conditioned by redshift, where each slice is flattened into a sequence of patches separated by newline tokens. Similar to a text-generation task, each pixel is treated as a token, and the model learns to predict the next token both spatially and across redshifts. However, such LLM-based approaches struggle to generalise beyond the dynamic range of their training data and often fail to generate coherent 3D 21-cm maps. A general limitation of these applications is that most map emulators are trained on computationally inexpensive simulations, such as small-scale 21cmFAST runs. Therefore, training these generative models from scratch typically requires  $10^5$ – $10^6$  samples, which is prohibitive for high-fidelity simulations. To address this, [Diao and Mao \(2025\)](#) proposed a pre-training and fine-tuning scheme in which a GAN pre-trained on small-scale simulations is fine-tuned to generate large-scale 21-cm maps using only  $\sim 10^2$  high-cost simulations.

The second class of AI applications to EoR simulations aims to accelerate or augment the modelling of physical processes that are too time-consuming or computationally expensive to resolve directly. These models are typically trained on high-fidelity simulations or constrained by galaxy observations, enabling them to capture sub-galactic-scale processes that would otherwise require prohibitively high resolution. One of the first examples was presented by [Sullivan et al. \(2017\)](#), who employed an artificial neural network to model radiative feedback in galaxies due to the UV background radiation. The network successfully connected the evolution of the baryon fraction and the suppression of star formation in galaxies during reionisation to local galaxy properties, including tidal forces, hydrogen ionisation fraction, gas temperature, and the halo virial ratio. Another example is provided by [Sherwin et al. \(2023\)](#), who used a fully connected neural network to link dark matter halo merger histories and galaxy properties with the UV escape fraction derived from high-redshift galaxy formation simulations. Deep neural networks have also been proposed to predict the stochastic star formation history of galaxies from dark matter halo properties at lower redshifts ( $z \sim 4$ ) (e.g. [Moster et al., 2021](#); [Behera et al., 2025](#)), and such approaches could be extended to higher redshifts. For instance, [Feathers et al. \(2025\)](#) employed a neural network as a sub-grid model to self-consistently determine the formation of the first stars in minihaloes ( $M_{\text{halo}} \leq 10^8 M_{\odot}$ ) in

the presence of a Lyman–Werner background during the early stages of reionisation. The network input consists of parameters from the surrounding large-scale environment of minihalos—such as the IGM overdensity and velocity field—and predicts the star formation rate in each simulation cell, significantly improving accuracy compared to semi-analytical models.

Moreover, the need for large cosmological volumes to reduce sample variance often limits the spatial resolution of simulations. Moreover, the minimum halo mass resolved is typically well above the expected threshold for haloes hosting star-forming galaxies. Here, AI-based super-resolution techniques play an increasingly important role, as they can be employed as alternatives to sub-grid models by artificially increasing the effective spatial or mass resolution of large-volume EoR simulations. The work of [Schaurecker et al. \(2022\)](#) employs a GAN architecture to populate low mass-resolution N-body simulations with additional smaller haloes, effectively lowering the minimum resolved halo mass by up to an order of magnitude in cGpc volumes N-body simulations. The resulting halo number counts agree well with theoretical predictions for a given cosmology. However, due to memory constraints, the network input consists of haloes being smoothed on increasingly fine grids rather than producing extended halo catalogues, meaning that the inpainted halo population remains distributed on a per-pixel level—thus limiting the achievable mass-resolution enhancement. Conversely, [Pochinda et al. \(2025\)](#) demonstrated that diffusion-based models can improve the resolution of cheap, low spatial-resolution simulations. Their score-based generative framework, combined with stochastic differential equations, can refine coarse but large volume 21-cm simulations to simultaneously reproduce the expected field of view and angular resolution of SKA-Low observations. While the approach recovers the 21-cm power spectrum with marginal errors below the anticipated SKA-Low noise level, it struggles to reproduce the sharp  $\delta T_b = 0$ , mK peak associated with ionised regions in the bi-modal 21-cm brightness distribution.

Overall, the integration of AI techniques into EoR simulations offers a promising pathway to bridge the gap between accuracy and computational feasibility. Neural networks provide a fast and cost-effective alternative for estimating global quantities, such as the volume-averaged neutral fraction or the timescale of reionisation, that are typically derived from high-fidelity but computationally expensive simulations. Although data-driven emulators and neural sub-grid models cannot yet fully replace radiative transfer and hydro-dynamical simulations, they have the potential to substantially accelerate parameter-space exploration and aid the calibration of numerical simulations. In the near future, combining carefully trained neural networks with traditional simulations may enable optimised inference from observations, ultimately enhancing our ability to interpret forthcoming 21-cm data from next-generation instruments such as SKA-Low.

### 3.2 Emulators of Summary Statistics

Observations can be robustly translated into constraints on astrophysics and cosmology via Bayesian inference (for more details see the Bayesian inference section in the Inference chapter; [Acharya et al. 2026](#)). State-of-the-art Bayesian inference requires multiple thousands of forward model evaluations (e.g.,  $\sim 200k$  in [HERA Collaboration et al. 2023](#)). Even the fastest semi-analytical models of the EoR (e.g., [Santos et al. 2010](#); [Hutter et al. 2021](#); [Davies et al. 2025](#)) take  $\sim 1h$  per evaluation, which can make such inferences computationally impractical. An emerging alternative to direct simulation has been the *emulation* of summary statistics. In contrast to the previous section

where one was emulating *fields* (e.g., the 21 cm brightness temperature field), in this section we discuss the use of surrogate models, often in the form of machine learning architectures like neural networks (NNs) to produce outputs such as power spectra from a simulator given a set of input parameters  $\theta$ . Emulators were first adopted within cosmology (e.g. see [Heitmann et al. 2006, 2009](#); [Fendt and Wandelt 2007](#); [Auld et al. 2007](#)) and have since been widely adopted in astrophysics (e.g. see [Kern et al. 2017](#); [Shimabukuro and Semelin 2017](#)). Emulation requires simulations but allows for *amortized* (i.e., reusable) inference. We still need to pay an upfront computational cost when we generate a suite of simulations to train the emulator (although online emulators have been developed in cosmology as well, see e.g. [Günther et al. 2025](#)). This set of simulations, however, is less computationally expensive than performing full inference with an expensive simulator, and provides full flexibility when running several inference rounds is required. The trained emulator can then be reused to quickly interpret observations as soon as they are released and test various theoretical models. Note that there is no *a priori* method to determine the number of training samples required to train a NN to a desired accuracy and precision. Recent work ([Lemos et al., 2024](#)) proposed a likelihood-free statistical test, PQMass, which quantifies whether two sets of samples—e.g., emulator outputs and simulations—are drawn from the same distribution, offering a principled way to assess generative model fidelity and sample sufficiency, and can therefore serve as a validation test of whether the training sample size is sufficient. As such, as instruments improve, more accurate and precise emulators will be required, which in turn shall become more and more expensive to train. To tackle this problem, the use of pre-trained models or foundational models which are trained on millions or even billions of observations is gaining popularity ([Parker et al., 2025b](#)), as reusing the information from previous runs can significantly lower the required computational budget.

Many current 21-cm instruments are focused on a detection of the spherically-averaged 1D 21-cm PS as it is a binned statistic that has enhanced signal-to-noise ratio in comparison to 3D maps. As such, the first emulators in the field naturally targeted the 1D PS (e.g. see [Kern et al. 2017](#); [Shimabukuro and Semelin 2017](#); [Schmit and Pritchard 2018](#); [Jennings et al. 2019](#); [Ghara et al. 2020](#); [Mondal et al. 2022](#); [Choudhury et al. 2025](#)). These models typically emulated the 1D PS to about 10% accuracy. After a claimed detection from EDGES of the globally-averaged 21-cm signal ([Bowman et al., 2018](#)), there has been a demand to emulate the global signal as well (e.g. see [Cohen et al. 2020](#); [Bevins et al. 2021](#); [Bye et al. 2022](#)). These emulators have sub-percent accuracy despite having generally increased the dimensionality of the parameter space compared to previous efforts. In some cases emulator efforts have incorporated more sophisticated NN architectures such as recursive neural networks (RNNs) ([Prelogović and Mesinger, 2023a](#)) and long-short-term memory (LSTM) ([Dorigo Jones et al., 2024](#)). With the ever-increasing sensitivity of the SKAO, the number of parameters to constrain will further increase, calling for novel techniques such as those relying on differentiable pipelines, which typically allow for a more efficient exploration of high-dimensional spaces.

The 21-cm signal is inherently non-Gaussian, as its evolution is governed by a complex interplay of highly non-linear astrophysical processes. Unlike the CMB, whose fluctuations are fully described by the power spectrum, the 21-cm signal contains rich higher-order correlations that encode additional information about the underlying physics. Consequently, the power spectrum—while

a convenient and widely used summary statistic—represents an imperfect compression of the full 21-cm field and inevitably discards valuable information. To address this limitation, recent studies have developed emulators targeting higher-order statistics such as the bispectrum, which are more sensitive to the non-linear structure of the signal.

Bispectrum computations are computationally challenging, especially when one is trying to compute all different shapes of unique triangles (Bharadwaj et al., 2020; Shaw et al., 2021), and when considering the large variation in dynamic range. A bispectrum emulator is therefore an attractive option for generating a plethora of bispectra samples in future Bayesian analysis of the signal. Recently, Tiwari et al. (2022) emulated the EoR 21cm bispectra using NNs, and used their emulator to forecast the astrophysical constraining power of bispectra.

One drawback to many emulators in the literature is that they produce only point value predictions of the target signal statistic; thus, they fail to capture the uncertainty in their predictions. Therefore, when such emulators are used in the Bayesian inference pipeline, they cannot naturally propagate their prediction uncertainties to the estimated model parameters. One way to propagate these errors is by using GPR to interpolate over errors that are cross-validated over a test dataset (Kern et al., 2017). The resulting emulator error covariances are then incorporated into one’s downstream analyses such as astrophysical parameter inference. A more sophisticated way is to develop an emulator using Bayesian neural networks (BNN), which can provide the uncertainty associated with its prediction as a model output. In Mahida et al. (2025), power spectrum and bispectrum emulators were developed using BNN architectures. These emulators provide the Gaussian posterior prediction of the predicted signal statistic in the form of the mean and covariance as model output. Thus, one can propagate this model prediction uncertainty to the Bayesian inference pipeline very easily.

## 4 Inference

Having discussed ML improvements to instrument modelling and data analysis in Sec. 2 and improvements to theoretical efforts in Sec. 3, we now bring the two together by considering the final step in any observational effort—the *inference* of how a particular astrophysical system (in our case, CD and EoR) behaves.

### 4.1 Auxiliary Data Products from Images

One application of ML models is to perform image segmentation to extract specific astrophysical features from tomographic data. This approach treats the map as an image to be classified pixel-by-pixel, enabling direct characterisation of the morphology and topology of structures during reionisation. U-shaped neural network architectures have been trained to produce binary maps that identify and delineate the morphology of HI and HII regions during the EoR (Bianco et al., 2021, 2024), even after brightness temperature maps have been distorted by foreground filtering (Gagnon-Hartman et al., 2021; Kennedy et al., 2024). This allows for a direct study of the topology of reionization—such as the size distribution, shape, and connectivity of ionized bubbles—a key science goal for the SKA-Low that can provide insights into the nature and distribution of the first ionizing sources (e.g., Giri et al., 2018a, 2019). Importantly, these segmentation

methods present a way to estimate the ionization history in a model-independent manner, as the volume of ionized or neutral regions can be directly measured from the identified features in the images, providing constraints on the reionization timeline without requiring assumptions about the underlying astrophysical processes.

Auxiliary data products arising from ML-processed images can also facilitate synergies with other probes of CD/EoR. For example, [Kennedy et al. \(2024\)](#) built on the work of [Gagnon-Hartman et al. \(2021\)](#) in recovering ionized bubbles in the presence of wedge-filtered 21 cm maps, generalizing the technique to operate on light cones and populating the simulated light cones with high-redshift galaxies. [Kennedy et al. \(2024\)](#) then showed that this ML-recovery of ionized regions can shrink the volume needed to look galaxies, potentially increasing the efficiency of high-redshift surveys. On the cosmological side, [Th  lie et al. \(2025\)](#) stacked the ML-recovered ionized bubbles to search for deviations from isotropy, enabling an Alcock-Paczyński test with the SKA to measure the product of the angular diameter distance and the Hubble parameter to  $\sim 2\%$  at reionization redshifts.

## 4.2 Parameter Inference

One important avenue where ML is transforming data-driven research in cosmology and astrophysics is through ML-based regression and inference ([Alsing et al., 2019](#); [Cranmer et al., 2020](#); [Spurio Mancini et al., 2022](#); [Villaescusa-Navarro et al., 2022](#); [Piras et al., 2024](#)). This links theory through model parameters directly to observations such as 21cm maps, i.e. without the need of an explicit likelihood. However, we often rely on pre-defined observables and low-dimensional summaries, limiting the use of rich observational data. The observed 21cm signal exhibits non-Gaussianity, non-analytic likelihoods, and complex structure, and is made even more complicated by systematics and foregrounds. We therefore profit from network-based solutions for optimal and unbiased inference that go beyond oversimplified and biased summary statistics such as power spectra. Moving towards field-level approaches has necessitated the development and adaptation of newer inference techniques such as amortized inference.

Besides constraining the astrophysics of galaxies and the intergalactic medium ([Park et al., 2019](#)), large-scale surveys such as mappings of the 21cm line with the SKAO enable studies of fundamental physics and cosmology ([Liu et al., 2020](#)). Fast simulation frameworks that are suitable for forward-modelling of synthetic observations and network training exist, for example, for generating 21cm 3D lightcones during cosmic dawn and reionization under different astrophysical ([Murray et al., 2020](#)) and cosmological ([Heneka and Amendola, 2018](#)) scenarios, complemented by smaller radiative hydrodynamical simulation datasets ([Meriot and Semelin, 2023](#)). A significant modeling error remains, as well as a lack of end-to-end forward modeling pipelines that include all noise systematics, requiring inference methods able to transfer well and provide robust posterior estimation. In addition, a framework to bridge the transfer gap between simulations and observations, as well as between different simulators is crucial.

For the 21cm signal, first inference applications of ML focused on regression of EoR parameters with dense neural networks (DNNs) from summaries such as 21cm power spectra ([Shimabukuro and Semelin, 2017](#)), later extended to extraction of summaries such as the global 21cm signal in the presence of foregrounds, ionospheric effects, instrumental responses, and thermal noise ([Choudhury](#)

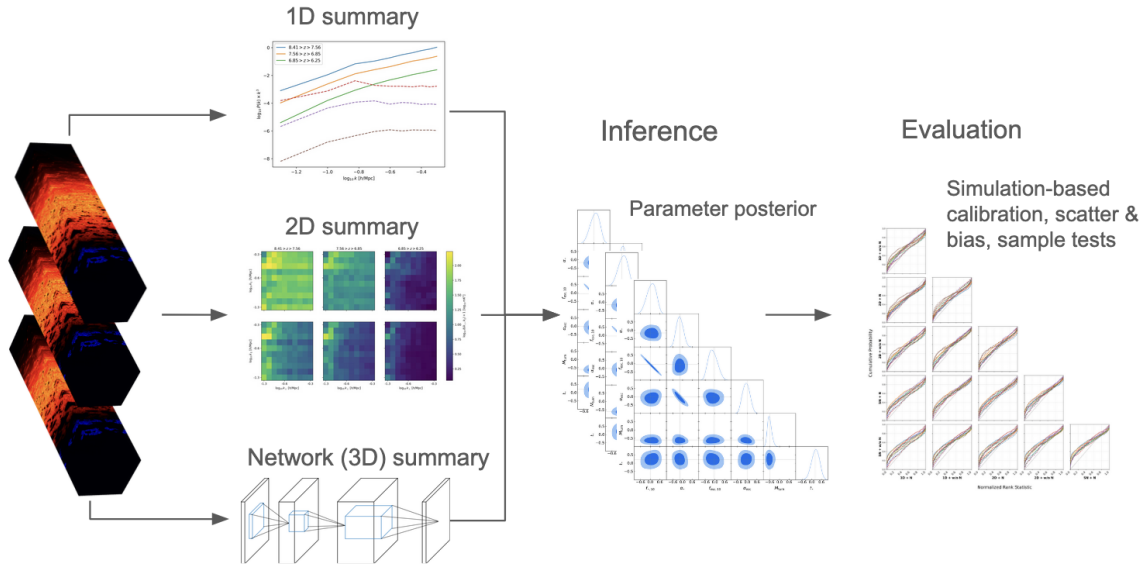
et al., 2020, 2021; Tripathi et al., 2024; Tripathi et al., 2024), followed by regression from power spectra (Choudhury et al., 2025) incorporating foreground effects (Choudhury et al., 2022). Since then, for parameter regression from full 21cm maps and light cones, convolutional neural networks (CNNs) have proven to work very well to summarise 21cm data for DNN regression of astrophysics and EoR properties (Gillet et al., 2019; Hassan et al., 2019; La Plante and Ntampaka, 2019; Kwon et al., 2020; Hassan et al., 2020; Mangena et al., 2020; Neutsch et al., 2022; Prelogović et al., 2022; Hiegel et al., 2023). Besides aleatoric uncertainty (inherent noise in the data) captured by the variance in regression point estimates, epistemic uncertainty (model-related) in relation to the chosen network-model can in part be captured with Bayesian Neural Networks (BNNs). However, these methods do not provide full posteriors on the inferred parameters. When coupled with MCMC or variational inference though, BNNs were employed to approximate posteriors of the 21cm signal (Hortúa et al., 2020; Meriot et al., 2025; Mahida et al., 2025).

For full recovery and direct modelling of posteriors, neural density estimators (NDEs) are used within simulation-based inference (SBI) frameworks, both for direct inference from 21cm maps and light cones and for inference via fixed summary statistics such as 21cm power spectra. Here, NDEs are employed in various tasks, including neural likelihood estimation (NLE), neural posterior estimation (NPE) and neural ratio estimation (NRE); they offer powerful approaches for fast, robust inference from high-dimensional data and for high-dimensional posteriors, without having to rely on explicit likelihoods or overly simplistic summary statistics (Lueckmann et al., 2018). As flexible network-based density estimators these methods typically utilize coupling layers or spline-based transformations within normalizing flows (Jimenez Rezende and Mohamed, 2015) or autoregressive models like masked autoregressive flows (Papamakarios et al., 2017). For example, NPE directly estimates the posterior distribution  $p(\theta|x)$  from simulated data, and has seen wide use for SBI with synthetic SKAO data, based on realistic SKA-Low antennae configurations such as AA4 and the inclusion of thermal noise, systematics, and foregrounds.

Given the complexity and high dimensionality of the 21cm data expected from the SKAO, the choice of suitable summary statistics is crucial for effective SBI. As power spectra require less integration time to achieve sensitivity for signal detection, they are expected to precede high-fidelity imaging. Therefore, SBI for fixed summary statistics derived from synthetic SKA-Low data has been studied in form of NDE based on the 21cm one-dimensional spherically averaged power spectrum (1DPS) and the two-dimensional cylindrically averaged power spectrum (2DPS), e.g. Saxena et al. (2023); Greig et al. (2024); Pietschke et al. (2025). As these studies have shown, SBI is able to derive EoR astrophysics and reconstruct the reionization history in an amortised and robust fashion. However, as for example Zhao et al. (2022b) and Prelogović and Mesinger (2023b) have found, when faced with synthetic data that includes noise and systematics, non-Gaussian estimators are needed even in case of the 1DPS to avoid overconfident and biased posteriors. Furthermore, first results from the SKA Science Data Challenge 3b on inference (SDC3b<sup>4</sup>), where the global ionization fraction at three different redshifts was to be blindly derived from 2DPS generated with two different simulators, point towards current posterior estimators to be susceptible to model misspecification both from signal and noise modelling; this has also been pointed out previously (Zhou and La Plante, 2022).

---

<sup>4</sup><https://sdc3.skao.int/challenges/inference>



**Figure 2:** Possible workflows for likelihood-free inference (Sec. 4): Once a database of 21 cm simulations is created, a suitable summary statistic (e.g. power spectra, higher-order statistics, or network-learned features) is extracted and used as input for simulation-based, likelihood-free inference based on neural density estimation. The resulting posteriors can be quantitatively evaluated using diagnostics such as simulation-based calibration and coverage tests.

Furthermore, the work of [Cerardi et al. \(2025\)](#) during the SDC3b shows that combining multiple summary statistics (2DPS, bispectrum or Betti numbers) can further enhance the extrapolation of global ionization fraction.

Finally, as an alternative to amortised inference with SBI based on a chosen summary statistic derived from simulated data, inference with for example MCMC becomes feasible when either leveraging ML for fast generative modelling of 21cm maps and light cones (see Sec. 3.1), or directly of statistics such as 21cm power spectra (see Sec. 3.2).

However, 1D and 2D power spectra are by construction Gaussian statistics and thus limited in fully characterising the inherently non-Gaussian 21cm signal; higher-order statistics such as the bispectrum can add constraining power but typically suffer from low S/N requiring accurate noise modelling ([Tiwari et al., 2022](#); [Mahida et al., 2025](#)). Alternatively, SBI offers the possibility to directly infer high-dimensional posteriors from 21cm maps and light cones, circumventing the need for a predefined statistic and instead employ neural networks to learn, also non-Gaussian, summary representations directly from the data. For the global 21cm signal, this has been done without a summarisation step ([Saxena et al., 2024](#)). In the case of high-dimensional 21cm maps and light cones as expected for SKA-Low, usually an SBI framework with NDE has been combined with a summarisation network, or compressor to first reduce dimensionality and learn a summary statistic, for example based on convolutional or information-maximizing neural network architectures ([Charnock et al., 2018](#)). This has been shown to provide a route to amortised, robust (unbiased) and optimal (neither over- nor underconfident) inference of astrophysical and cosmological information

from synthetic SKA data (Prelogović and Mesinger, 2024; Schosser et al., 2025a,b). In particular, data summarisation based on 3D convolutions has been shown to be efficient when dealing with 21cm light cones (Neutsch et al., 2022; Zhao et al., 2022a). As has been noted in these studies, a joint optimisation of the network-learnt summary and posterior estimation is crucial in order to ensure robust and optimal parameter estimates. In addition, hybrid summaries and wavelets have emerged as an approach to combine physically motivated statistics with network-learned features to leverage both interpretability and data-driven optimization (Zhao et al., 2023, 2024; Makinen et al., 2024).

Figure 2 depicts a workflow for ML-based inference of EoR model posteriors from SKA-Low data. We first derive either a fixed summary, such as a 1DPS or 2DPS, or a network-learnt summary as a lower-dimensional representation of high-dimensional SKA maps and light cones. As described in this section, within SBI frameworks for likelihood-free inference NDEs trained on simulated data directly estimate the high-dimensional parameter posterior. Also, fast emulators of simulations and summaries as described in Sec. 3.1 and Sec. 3.2 can be used for either likelihood-based Bayesian sampling or coupled to likelihood-free inference. Finally, accuracy and precision of derived posteriors can be quantitatively evaluated. For regression tasks, scatter plots that compare mean and variance for many realisations to determine bias are common. Quantitative approaches for inference include simulation-based calibration (SBC; Talts et al., 2018), posterior coverage tests (Cook et al., 2006), and Tests of Accuracy with Random Points (TARP; Lemos et al., 2023). The majority of these methods are in principle also applicable to likelihood-based inference. Due to the intractable likelihood of the 21cm signal, for the SKA their use is enabled by either fast joint generation of parameters and signal data, or amortised posterior evaluation speed for SBI as compared to, for example, MCMC based on sampling 21cm simulations. One or more of the above-mentioned methods therefore are routinely used to judge posterior performance in 21cm inference with SBI.

In a next step, studies of likelihood-free 21cm inference have recently explored routes to improved network-learnt summaries that are more robust than standard supervised approaches. The main motivation is an inherent problem for 21cm inference that arises due to differences observed across different 21cm forward models (model misspecification). Even when employing identical physical assumptions, distinct simulators can produce notably different outputs, introducing potential biases for an inference model trained on samples from only one simulator (Sooknunan et al., 2024). Related approaches such as Evidence Networks (Jeffrey and Wandelt, 2024) demonstrate that neural estimators can perform amortized, simulation-based model comparison, offering a way to test and mitigate biases arising from mismatched simulators; and Diao et al. (2025a) shows that neural density estimators such as continuous-time flow models can diagnose and quantify biases between forward models. Beyond discrepancies between simulators, inference methods must also handle variations in observational conditions, including instrumental noise models, resolutions, and systematics, factors crucially tested, for example, by SKA SDC3b. Options so far explored for 21cm summaries are information-maximising networks (Prelogović and Mesinger, 2024); whether they are maximally informative given the data is evaluated by comparing posterior variance against Fisher information or the Cramér–Rao bound. Besides augmentation of training data, semi-supervised and self-supervised learning techniques based on masking and contrastive learning strategies,

e.g., [LeCun \(2022\)](#), also provide a promising route toward addressing these robustness concerns. Indeed, self-supervised vision transformers trained on low-resolution 21cm light cones have been shown to generalize well to unseen higher-resolution, noised data and new model parameters ([Ore et al., 2025](#)), on the way to potentially exploit industry-sized models ([Heneka et al., 2025](#)).

Finally, it is very likely that synergies with other observables will ultimately be required to obtain the tightest constraints on our observables and to learn the most about CD/EoR. State-of-the-art inferences therefore combine with other probes such as UV luminosity functions that allow for the inclusion of JWST and HST observations, Thomson scattering optical depth to the CMB that allows for the inclusion of Planck observations, as well as full multi-probe analysis for example with CMB power spectra. A range of studies has already used ML for multi-probe joint inference with the SKA, precursor instruments such as HERA, and 21cm global signal measurements, using successfully NDE to evaluate joint likelihoods ([Bevins et al., 2023, 2024](#); [Sims et al., 2025](#)), or for joint NPE with conditional flow matching ([Schosser et al., 2025b](#)); these studies have shown that SBI can robustly combine other probes with the 21cm signal in a likelihood-free way, leading to improved constraints on CD/EoR physics.

## 5 Some Developing Trends

In this chapter, we have discussed a number of advances in the application of ML to CD/EoR science. Given the infancy of this flavour of research, it would be fair to say that the role of ML in the SKA remains speculative for now, and may evolve considerably in the next few years.

However, it would be reasonable to guess that one low-level technique is likely to play a role: automatic differentiation (AD). AD is an emerging approach to solving difficult, high dimensional inverse problems, where one can compute the gradient of an objective function automatically with respect to its model parameters. These gradients can then be used for optimization or sampling, and are particularly useful when working in high-dimensional parameter spaces. Methods for computing gradients of a data model include: 1) numerical finite-difference approach, which is simple but contains numerical noise and can be prohibitively expensive in high dimensions; 2) symbolic regression, which computes exact gradients but is confined to working with closed-form analytic expressions, and 3) AD, which uses the chain rule of differential calculus to compute exact gradients of any series of numerical operations. With access to gradients, one can employ a wide range of gradient-aware optimization techniques, such as 1st-order gradient descent algorithms and their derivatives (e.g. SGD, Adam), and 2nd-order gradient descent algorithms and their derivatives (e.g. BFGS, L-BFGS). Furthermore, sampling the objective function, as is done with Markov Chain Monte Carlo approaches, can be made more efficient with gradients via algorithms like Hamiltonian Monte Carlo and its derivatives (e.g. NUTS).

If GPU-compute and large-scale datasets are the fuel of the AI revolution, then AD can be thought of as the train powering that revolution. Indeed, it was the advent of flexible, user-friendly AD libraries that enabled the widespread adoption of deep learning ([Gunes Baydin et al., 2015](#), and references therein). While AD powers the training of neural networks, it can also be used to accelerate and scale-up our ability to solve traditional inverse problems, without the use of neural networks. This general

paradigm of programming with AD-enabled libraries is known as *differentiable programming*, which has seen widespread adoption within the physical sciences. The benefit here is the ability to leverage the inductive bias of our physics-based forward models, while still reaping the benefit of gradient-aware optimization and easy GPU portability afforded by high-level AD libraries (e.g. PyTorch, Tensorflow, JAX, Julia).

In 21 cm cosmology, this has only recently begun to take hold. [Yatawatta et al. \(2019\)](#) presented an approach for interferometric calibration using an L-BFGS optimizer coupled with an AD-enabled calibration forward model, demonstrating improved performance over standard calibration techniques. [Ewall-Wice et al. \(2022\)](#) presented a framework for the joint modeling of foregrounds and antenna beams in interferometric visibilities based on the principle of minimizing spectral structure. This allows for the subtraction of beam-weighted foregrounds with minimal a priori assumptions about the structure of the beam. Their forward model was wrapped in an AD-enabled library, and thus allowed for straightforward gradient-descent optimization and GPU-acceleration. [Cox et al. \(2024\)](#) took this a step further, showing how to leverage spectral redundancy in such an approach and thus improve the quality of antenna gain calibration. [Diao et al. \(2025b\)](#) introduced a differentiable simulation for synchrotron from Galactic fields, which can be used for extract information from observation or integrated into AD-enabled pipelines.

Recently, [Kern \(2025\)](#) proposed an *end-to-end* differentiable forward model for 21 cm cosmological analysis. In their forward model, signals are parameterized in their native space (e.g. on the sky), and then forward modeled to the interferometric visibilities via the measurement equation ([Smirnov, 2011](#)). Wrapping this with an AD-enabled library allows one to differentiate backwards through the measurement process, and derive gradients of a visibility-based loss function with respect to sky pixels. Coupled with an L-BFGS solver allows for joint optimization of foreground, instrument, and 21 cm signal parameters. Additionally, coupled with a Hamiltonian Monte Carlo algorithm ([Neal, 2011](#)) allows for joint posterior density estimation and nuisance parameter marginalization. This forward model is shown in [Figure 3](#). Differentiable programming is crucial to making this feasible on large-scale datasets. ([Kern, 2025](#)) present a proof-of-concept for a scaled-down version of the HERA telescope, showing robust joint posterior density estimation between diffuse and point source foregrounds, a complex frequency and angle dependent antenna beam model, and the 21 cm sky signal, with a total walltime of 16 hours spread across 4 GPUs. Going forward, AD-enabled pipelines will help a variety of tasks within the standard 21 cm workflow, such as antenna calibration, beam modeling, and foreground subtraction.

Another important future prospect that cannot be ignored is the proliferation of modern AI tools based on foundation models, most notably large language models (LLMs). LLMs have achieved remarkable success in coding, summarizing papers, and synthesizing related literature. These tools are already being integrated into the daily workflows of many researchers, including many of those within the SKA 21-cm community. Looking ahead, foundation models will likely shape the paradigm of our research in two distinct ways: explicitly and implicitly.

Explicitly, foundation models can be trained to directly perform scientific tasks, including but not limited to map generation, parameter estimation, and source identification. There are two primary approaches to deploying such models. The first is training from scratch. There have been successful

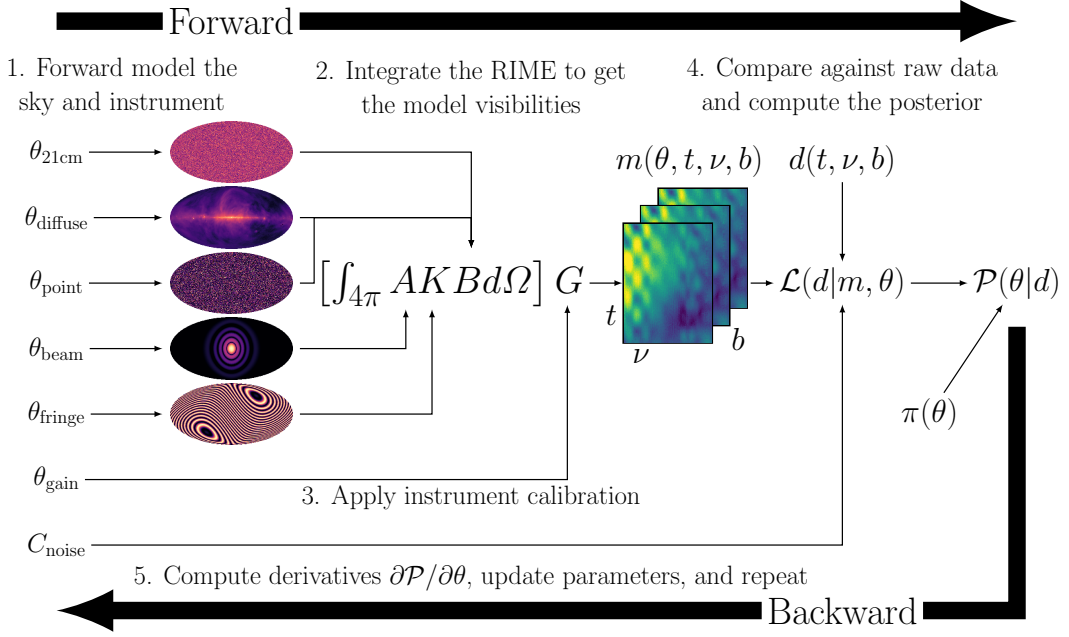
cases in astrophysics, particularly in fields where the number of catalogued astronomical objects reaches the hundreds of millions (e.g. Parker et al., 2025a; Zhao et al., 2026; Parker et al., 2024), enabling powerful performance in tasks like redshift and stellar parameter estimation. However, aggregating comparable volumes of training data remains highly challenging for 21 cm line intensity mapping and large-scale structure studies. Instead of training from scratch, the second approach is to fine-tune an existing foundation model. Fine-tuning allows one to retrain a general task foundation model on small dataset with limited training cost but achieve stronger performances on specific tasks (see e.g. Balavadhani Parthasarathy et al. (2024) for a review). Because foundation models possess generalized capabilities learned across a wide range of domains, fine-tuning them with a relatively small amount of domain-specific data is often sufficient to achieve high performance on specialized tasks. For example, the work of Heneka et al. (2025) fine-tuned existing foundation models, such as Qwen, on 21cmFAST simulations to analyze and generate 21-cm tomographic maps from astrophysical and cosmological parameters, confirming the viability of repurposing general-purpose foundation models for specialized cosmological applications. A true SKA-foundation model that can transfer between, e.g. training domains for different SKA-related tasks, similar to what has been done in Sec. 4, remains to be built and tested, representing an interesting challenge for researchers to pursue.

Beyond these explicit applications, foundation models will also contribute to scientific output implicitly. By leveraging these tools, researchers can more effectively communicate scientific progress in papers and presentations, while AI-generated literature summaries will make complex research more accessible. Furthermore, the community will greatly benefit from AI coding assistants, which promise to streamline software development and data analysis pipelines. Recently, there are even trends of AI agents (e.g. Villaescusa-Navarro et al., 2025), constructing an automatic workflow of all above procedures with possible human interaction, accelerating the overall pace of scientific discovery.

## 6 Summary and Future Prospects

In this chapter, we have examined applications of machine learning across the full analysis pipeline of CD/EoR science—from the observational data analysis to its theoretical interpretation. The breadth and depth of these applications suggest that ML will be a serious contender in the toolkit for realising the full scientific potential of the SKA in the long term.

The methods described in Section 2 demonstrate a fundamental shift in how we approach observational data analysis. Traditional approaches to RFI mitigation, ionospheric correction, and gain calibration have relied on carefully engineered heuristics and statistical outlier detection. ML methods, on the other hand, can learn complex, non-linear patterns directly from data, adapting to the specific characteristics of each observation. For example, recent results from LOFAR demonstrate that improved gain calibration coupled with ML-based foreground mitigation can significantly improve data quality. As SKA-Low begins operations, we anticipate that similar methods will be essential and might transition from research prototypes as they currently are to operational components of standard analysis pipelines. The key challenge with such changes will be ensuring robustness and interpretability: ML models must not introduce biases or spurious features that could



**Figure 3:** A differentiable, Bayesian forward model for 21 cm cosmology from Kern (2025), highlighting the *forward* operation of visibility simulation, and the *backward* operation of reverse-mode backpropagation. Optimization and sampling of the (un-normalized) posterior distribution proceeds iteratively, using the computed gradients to move within the joint parameter space of the model.

be mistaken for cosmological signal. Techniques such as explainable AI, uncertainty quantification, and rigorous validation against simulated datasets will be crucial in the coming years to build trust in these methods.

These advances in observational pipelines are complemented by equally significant progress on the theoretical side. The computational demands of EoR simulations represent a fundamental bottleneck to interpreting observations with forward modelling within an inference context. As discussed in Section 3, field-level and summary statistic emulators offer a path forward, enabling inference at scales that would otherwise be computationally prohibitive. The next generation of emulators could incorporate several key advances: multi-fidelity approaches that combine cheap low-resolution simulations with expensive high-resolution runs to emulate large simulation boxes at high resolution, and physics-informed architectures that embed known symmetries and conservation laws. As emulators become more accurate and widely adopted, we may see a convergence between simulation-based and analytic modeling approaches, with differentiable simulators enabling end-to-end gradient-based inference through the full forward model.

Bridging theory and observation, Section 4 highlighted the shift toward likelihood-free, simulation-based inference methods that can extract information directly at the field level. This represents a shift away from a set of pre-defined summary statistics toward more optimal, data-driven approaches. Field-level inference will likely become routine, with NNs trained to map directly from observed data cubes to posterior distributions over astrophysical parameters. Hierarchical inference frameworks will enable simultaneous characterization of astrophysical populations and instrumental

systematics, avoiding the artificial separation between “signal” and “nuisance” parameters. Multi-probe inferences will combine 21-cm data with complementary observations (e.g., see Synergies and Inference chapters; [Chakraborty et al. 2026](#); [Acharya et al. 2026](#)) using ML methods designed for such synergistic analyses. A critical challenge will be ensuring that these methods produce well-calibrated uncertainties, particularly in the presence of model misspecification (see Inference chapter for more discussion on this; [Acharya et al. 2026](#)).

The development of autodifferentiable frameworks, as outlined in Section 5, provides the technical foundation for many of these advances. Automatic differentiation enables gradient-based optimization and sampling in high-dimensional spaces, and the coming years will likely see the development of fully differentiable end-to-end pipelines that connect raw visibilities directly to astrophysical parameters, with gradients flowing through instrumental models, foreground components, and cosmological forward models. This will enable joint optimization over all aspects of the analysis.

While this chapter discusses ML applications to the individual parts of the CD/EoR science pipeline, the greatest impact may come from combining these methods. For example, a single differentiable pipeline might simultaneously perform RFI flagging, foreground removal, and parameter inference, with each component using some of the ML methods described in the corresponding subsections. Such end-to-end approaches would lead to extracting the maximal amount of information from upcoming observations.

A growing trend in machine learning is the development of foundation models: large, general-purpose models trained on diverse datasets that can be fine-tuned for specific tasks with relatively little additional data. In astronomy, this paradigm is also slowly beginning to take hold, with models pre-trained on large simulation suites or archival observations being adapted to new instruments or science cases. For CD/EoR science, foundation models could enable transfer learning from related domains such as galaxy surveys to SKA observations, potentially offering new ways to combine 21-cm data with complementary information.

Despite the impressive progress made so far, significant challenges that cannot be ignored remain, especially in terms of explainability and robustness. ML methods are often treated as “black boxes,” making them difficult to understand, which is problematic for science in general, but becomes especially problematic when those methods fail. Ensuring that ML-driven methods are robust against adversarial inputs, out-of-distribution data, and model misspecification will require significant ongoing methodological development.

The SKA is poised to be one of the defining radio observatories of the next decade. With its high sensitivity and high data rates, it will almost certainly be a key instrument in taking advantage of what ML has to offer to astrophysics. The challenge for the community in the coming years until then will be to develop, validate, and deploy ML methods responsibly, ensuring that they enhance rather than obscure our understanding of the Universe’s first billion years.

## 7 Author Contributions

A. Liu, C. Heneka, and M. Bianco coordinated the writing of this chapter, organised the contributions and edited the content. They also provided a final general overview throughout the chapter, improving readability and formatting. We list the author’s contributions for each section below.

**Sec. 1** A. Liu wrote this section.

**Sec. 2** C. Sui wrote the Sec. 2.1. S. K. Pal, A. Tripathi and A. Datta wrote Sec. 2.2 on ionospheric effects and provided a first draft of Sec. 2.3 on gain calibration. S. K. Giri wrote the first paragraph of the imaging section, Sec. 2.4, while Shulei Ni wrote the remaining paragraphs. A. Acharya primarily wrote Sec. 2.5 with contributions from S. K. Giri and D. Piras.

**Sec. 3** M. Bianco primarily wrote Sec. 3.1. Y. Mahida and S. Majumdar implemented the paragraph about the CosmoUiT, while C. Heneka implemented the LLMs part, and N. Kern wrote key details on emulation challenges and historical applications. S. K. Giri, D. Breitman and D. Piras provided general contributions to the section. D. Breitman coordinated the writing of Sec. 3.2, while A. K. Shaw, A. Tripathi, Y. Mahida and S. Majumdar primarily contributed to the fourth and fifth paragraphs. K. Diao contributed to the introduction of the map-level emulators in Sec. 3.1, focusing on GAN, diffusion, and model fine-tuning. X. Zhao contributed to the editing of Sec. 3.

**Sec. 4** A. Liu wrote Sec. 4.1. Y. Pietschke and C. Heneka primarily wrote Sec. 4.2. A. Tripathi and Y. Mahida wrote the segments on the global signal and the application of BNN. X. Zhao contributed to the editing of Sec. 4. H. Shimabukuro contributed to the third paragraph at Sec 4.2.

**Sec. 5** A. Liu, N. Kern, and K. Diao wrote this section.

**Sec. 6** D. Breitman wrote this section, with minor edits by A. Liu.

## References

- A. Acharya et al. *MNRAS*, 527(3):7835–7846, Jan. 2024a. doi: 10.1093/mnras/stad3701.
- A. Acharya et al. *MNRAS*, 534(1):L30–L34, Oct. 2024b. doi: 10.1093/mnras/slae078.
- A. Acharya et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Acharya01.
- J. Akeret, C. Chang, A. Lucchi, and A. Refregier. *Astronomy and Computing*, 18:35–39, Jan. 2017. doi: 10.1016/j.ascom.2017.01.002.
- J. G. Albert, R. J. van Weeren, H. T. Intema, and H. J. A. Röttgering. *A&A*, 635:A147, Mar. 2020. doi: 10.1051/0004-6361/201937424.
- J. Alsing, T. Charnock, S. Feeney, and B. Wandelt. *MNRAS*, 488(3):4440–4458, Sept. 2019. doi: 10.1093/mnras/stz1960.
- D. Aubert and R. Teyssier. *The Astrophysical Journal*, 724(1):244, nov 2010. doi: 10.1088/0004-637X/724/1/244. URL <https://dx.doi.org/10.1088/0004-637X/724/1/244>.

- D. Aubert, N. Deparis, and P. Ocvirk. *Monthly Notices of the Royal Astronomical Society*, 454(1):1012–1037, 09 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv1896. URL <https://doi.org/10.1093/mnras/stv1896>.
- T. Auld, M. Bridges, M. P. Hobson, and S. F. Gull. *MNRAS*, 376(1):L11–L15, Mar. 2007. doi: 10.1111/j.1745-3933.2006.00276.x.
- W. A. Baan, P. A. Fridman, and R. P. Millenaar. *AJ*, 128(2):933–949, Aug. 2004. doi: 10.1086/422350.
- V. Balavadhani Parthasarathy, A. Zafar, A. Khan, and A. Shahid. *arXiv*, art. arXiv:2408.13296, Aug. 2024. doi: 10.48550/arXiv.2408.13296.
- P. Baldi. In I. Guyon et al., editors, *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, volume 27 of *Proceedings of Machine Learning Research*, pages 37–49, Bellevue, Washington, USA, 02 Jul 2012. PMLR. URL <https://proceedings.mlr.press/v27/baldi12a.html>.
- N. Barry et al. *MNRAS*, 461(3):3135–3144, Sept. 2016. doi: 10.1093/mnras/stw1380.
- A. P. Beardsley et al. *ApJ*, 833(1):102, Dec. 2016. doi: 10.3847/1538-4357/833/1/102.
- J. Behera, R. Tojeiro, and H. G. Chittenden. Optimised neural network predictions of galaxy formation histories using semi-stochastic corrections, 2025. URL <https://arxiv.org/abs/2409.16548>.
- E. Beohar et al. Mitigating gain calibration errors from eor observations with ska1-low aa\*, 2025. URL <https://arxiv.org/abs/2510.25886>.
- H. T. J. Bevins et al. *MNRAS*, 508(2):2923–2936, Dec. 2021. doi: 10.1093/mnras/stab2737.
- H. T. J. Bevins et al. *MNRAS*, 526(3):4613–4626, Dec. 2023. doi: 10.1093/mnras/stad2997.
- H. T. J. Bevins et al. *MNRAS*, 527(1):813–827, Jan. 2024. doi: 10.1093/mnras/stad3194.
- S. Bharadwaj, A. Mazumdar, and D. Sarkar. *MNRAS*, 493(1):594–602, Mar. 2020. doi: 10.1093/mnras/staa279.
- M. Bianco, S. K. Giri, I. T. Iliev, and G. Mellema. *Monthly Notices of the Royal Astronomical Society*, 505(3):3982–3997, May 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab1518. URL <http://dx.doi.org/10.1093/mnras/stab1518>.
- M. Bianco et al. *Monthly Notices of the Royal Astronomical Society*, 528(3):5212–5230, Jan. 2024. ISSN 1365-2966. doi: 10.1093/mnras/stae257. URL <http://dx.doi.org/10.1093/mnras/stae257>.
- M. Bianco et al. *Monthly Notices of the Royal Astronomical Society*, 541(1):234–250, June 2025. ISSN 1365-2966. doi: 10.1093/mnras/staf973. URL <http://dx.doi.org/10.1093/mnras/staf973>.

- A. Bonaldi et al. *arXiv e-prints*, art. arXiv:2503.11740, Mar. 2025. doi: 10.48550/arXiv.2503.11740.
- J. D. Bowman et al. *Nature*, 555(7694):67–70, Mar. 2018. doi: 10.1038/nature25792.
- S. A. Brackenhoff et al. *MNRAS*, 533(1):632–656, Sept. 2024. doi: 10.1093/mnras/stae1856.
- P. R. Burd et al. *Astronomische Nachrichten*, 339(5):358–362, June 2018. doi: 10.1002/asna.201813505.
- C. H. Bye, S. K. N. Portillo, and A. Fialkov. *ApJ*, 930(1):79, May 2022. doi: 10.3847/1538-4357/ac6424.
- R. Byrne, M. F. Morales, B. J. Hazelton, and M. Wilensky. *MNRAS*, 503(2):2457–2477, May 2021. doi: 10.1093/mnras/stab647.
- E. Ceccotti et al. *arXiv e-prints*, art. arXiv:2504.18534, Apr. 2025. doi: 10.48550/arXiv.2504.18534.
- N. Cerardi et al. *arXiv preprint arXiv:2510.14877*, 2025.
- A. Chakraborty et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Chakraborty01.
- J. Chardin et al. *Monthly Notices of the Royal Astronomical Society*, 490(1):1055–1065, 09 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz2605. URL <https://doi.org/10.1093/mnras/stz2605>.
- T. Charnock, G. Lavaux, and B. D. Wandelt. *Phys. Rev. D*, 97(8):083004, Apr. 2018. doi: 10.1103/PhysRevD.97.083004.
- S.-F. Chen, K.-F. Chen, and C. Dvorkin. Field-level reconstruction from foreground-contaminated 21-cm maps, 2025. URL <https://arxiv.org/abs/2508.13265>.
- T. Chen et al. *MNRAS*, 532(2):2615–2634, Aug. 2024. doi: 10.1093/mnras/stae1676.
- Z. Chen, E. Chapman, L. Wolz, and A. Mazumder. *MNRAS*, 524(3):3724–3740, Sept. 2023. doi: 10.1093/mnras/stad2102.
- M. Choudhury, A. Datta, and A. Chakraborty. *MNRAS*, 491(3):4031–4044, Jan. 2020. doi: 10.1093/mnras/stz3107.
- M. Choudhury, A. Chatterjee, A. Datta, and T. R. Choudhury. *MNRAS*, 502(2):2815–2825, Apr. 2021. doi: 10.1093/mnras/stab180.
- M. Choudhury, A. Datta, and S. Majumdar. *MNRAS*, 512(4):5010–5022, June 2022. doi: 10.1093/mnras/stac736.
- M. Choudhury et al. *JCAP*, 2025(6):003, June 2025. doi: 10.1088/1475-7516/2025/06/003.
- A. Cohen, A. Fialkov, R. Barkana, and R. A. Monsalve. *MNRAS*, 495(4):4845–4859, July 2020. doi: 10.1093/mnras/staa1530.

- S. R. Cook, A. Gelman, and D. B. Rubin. *Journal of Computational and Graphical Statistics*, 15 (3):675–692, 2006.
- T. A. Cox et al. *Mon. Not. Roy. Astron. Soc.*, 532(3):3375–3394, 2024. doi: 10.1093/mnras/stae1612.
- K. Cranmer, J. Brehmer, and G. Louppe. *Proceedings of the National Academy of Science*, 117 (48):30055–30062, Dec. 2020. doi: 10.1073/pnas.1912789117.
- A. Datta, S. Bhatnagar, and C. L. Carilli. *ApJ*, 703(2):1851–1862, Oct. 2009. doi: 10.1088/0004-637X/703/2/1851.
- A. Datta et al. *arXiv e-prints*, art. arXiv:1409.0513, Sept. 2014. doi: 10.48550/arXiv.1409.0513.
- J. E. Davies, A. Mesinger, and S. Murray. *arXiv e-prints*, art. arXiv:2504.17254, Apr. 2025. doi: 10.48550/arXiv.2504.17254.
- J. M. G. H. J. de Jong et al. *MNRAS*, 542(4):3253–3276, Oct. 2025. doi: 10.1093/mnras/staf1373.
- K. Diao and Y. Mao. *The Astrophysical Journal*, 992(2):173, oct 2025. doi: 10.3847/1538-4357/ae0325. URL <https://doi.org/10.3847/1538-4357/ae0325>.
- K. Diao, B. Dai, and U. Seljak. *JCAP*, 2025(8):004, Aug. 2025a. doi: 10.1088/1475-7516/2025/08/004.
- K. Diao, Z. Li, R. D. P. Grumitt, and Y. Mao. *ApJS*, 278(1):25, May 2025b. doi: 10.3847/1538-4365/adc5ff.
- J. Dorigo Jones et al. *ApJ*, 977(1):19, Dec. 2024. doi: 10.3847/1538-4357/ad8b20.
- C. D. Du Toit, T. L. Grobler, and D. J. Ludick. *MNRAS*, 530(1):613–629, May 2024. doi: 10.1093/mnras/stae892.
- C. Dvorkin et al. Machine learning and cosmology, 2022. URL <https://arxiv.org/abs/2203.08056>.
- A. Ewall-Wice et al. *ApJ*, 938(2):151, Oct. 2022. doi: 10.3847/1538-4357/ac87b3.
- C. R. Feathers, M. Kulkarni, and E. Visbal. *Journal of Cosmology and Astroparticle Physics*, 2025(02):043, Feb. 2025. ISSN 1475-7516. doi: 10.1088/1475-7516/2025/02/043. URL <http://dx.doi.org/10.1088/1475-7516/2025/02/043>.
- W. A. Fendt and B. D. Wandelt. *arXiv e-prints*, art. arXiv:0712.0194, Dec. 2007. doi: 10.48550/arXiv.0712.0194.
- S. Gagnon-Hartman, Y. Cui, A. Liu, and S. Ravanbakhsh. *Monthly Notices of the Royal Astronomical Society*, 504(4):4716–4729, Apr. 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab1158. URL <http://dx.doi.org/10.1093/mnras/stab1158>.
- E. Garaldi et al. *MNRAS*, 512(4):4909–4933, June 2022. doi: 10.1093/mnras/stac257.
- B. K. Gehlot et al. *MNRAS*, 488(3):4271–4287, Sept. 2019. doi: 10.1093/mnras/stz1937.
- F. Geyer et al. *A&A*, 677:A167, Sept. 2023. doi: 10.1051/0004-6361/202347073.

- R. Ghara, T. R. Choudhury, and K. K. Datta. *MNRAS*, 447(2):1806–1825, Feb. 2015. doi: 10.1093/mnras/stu2512.
- R. Ghara et al. *Monthly Notices of the Royal Astronomical Society*, 476(2):1741–1755, Feb. 2018. ISSN 1365-2966. doi: 10.1093/mnras/sty314. URL <http://dx.doi.org/10.1093/mnras/sty314>.
- R. Ghara et al. *MNRAS*, 493(4):4728–4747, Feb. 2020. doi: 10.1093/mnras/staa487.
- N. Gillet et al. *MNRAS*, 484(1):282–293, Mar. 2019. doi: 10.1093/mnras/stz010.
- S. K. Giri, G. Mellema, K. L. Dixon, and I. T. Iliev. *Monthly Notices of the Royal Astronomical Society*, 473(3):2949–2964, 2018a.
- S. K. Giri, G. Mellema, and R. Ghara. *Monthly Notices of the Royal Astronomical Society*, 479(4):5596–5611, 2018b.
- S. K. Giri et al. *Monthly Notices of the Royal Astronomical Society*, 489(2):1590–1605, 2019.
- S. K. Giri et al. *Monthly Notices of the Royal Astronomical Society*, 533(2):2364–2378, 08 2024. ISSN 0035-8711. doi: 10.1093/mnras/stae1999. URL <https://doi.org/10.1093/mnras/stae1999>.
- I. J. Goodfellow et al. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, page 2672–2680, Cambridge, MA, USA, 2014. MIT Press.
- B. Greig et al. *MNRAS*, 533(2):2530–2545, Sept. 2024. doi: 10.1093/mnras/stae1984.
- A. Gunes Baydin, B. A. Pearlmutter, A. Andreyevich Radul, and J. M. Siskind. *arXiv e-prints*, art. arXiv:1502.05767, Feb. 2015. doi: 10.48550/arXiv.1502.05767.
- S. Günther et al. *JCAP*, 2025(9):059, Sept. 2025. doi: 10.1088/1475-7516/2025/09/059.
- S. Hassan, A. Liu, S. Kohn, and P. La Plante. *MNRAS*, 483(2):2524–2537, Feb. 2019. doi: 10.1093/mnras/sty3282.
- S. Hassan, S. Andrianomena, and C. Doughty. *MNRAS*, 494(4):5761–5774, June 2020. doi: 10.1093/mnras/staa1151.
- K. Heitmann, D. Higdon, C. Nakhleh, and S. Habib. *ApJL*, 646(1):L1–L4, July 2006. doi: 10.1086/506448.
- K. Heitmann et al. *ApJ*, 705(1):156–174, Nov. 2009. doi: 10.1088/0004-637X/705/1/156.
- C. Heneka and L. Amendola. *JCAP*, 2018(10):004, Oct. 2018. doi: 10.1088/1475-7516/2018/10/004.
- C. Heneka et al. *arXiv e-prints*, art. arXiv:2506.14757, June 2025. doi: 10.48550/arXiv.2506.14757.
- HERA Collaboration et al. *ApJ*, 945(2):124, Mar. 2023. doi: 10.3847/1538-4357/acaf50.
- J. Hiegel et al. *A&A*, 679:A125, Nov. 2023. doi: 10.1051/0004-6361/202346495.

- P. Hirling et al. *Astronomy and Computing*, 48:100861, July 2024. ISSN 2213-1337. doi: 10.1016/j.ascom.2024.100861. URL <http://dx.doi.org/10.1016/j.ascom.2024.100861>.
- J. Ho, A. Jain, and P. Abbeel. In H. Larochelle et al., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf).
- J. Högbom. *Astronomy and Astrophysics Supplement*, Vol. 15, p. 417, 15:417, 1974.
- H. J. Hortúa, R. Volpi, and L. Malagò. *arXiv e-prints*, art. arXiv:2005.02299, May 2020. doi: 10.48550/arXiv.2005.02299.
- I. Hothi et al. *MNRAS*, 500(2):2264–2277, Jan. 2021. doi: 10.1093/mnras/staa3446.
- I. Hothi, E. Allys, B. Semelin, and R. Meriot. *arXiv e-prints*, art. arXiv:2506.10061, June 2025. doi: 10.48550/arXiv.2506.10061.
- I. Hothi et al. *Mon. Not. Roy. Astron. Soc.*, 500(2):2264–2277, 2020. doi: 10.1093/mnras/staa3446.
- A. Hutter et al. *MNRAS*, 503(3):3698–3723, May 2021. doi: 10.1093/mnras/stab602.
- I. T. Iliev et al. *Monthly Notices of the Royal Astronomical Society*, 439(1):725–743, Jan. 2014. ISSN 0035-8711. doi: 10.1093/mnras/stt2497. URL <http://dx.doi.org/10.1093/mnras/stt2497>.
- B. Jacobson-Bell et al. *AJ*, 169(4):206, Apr. 2025. doi: 10.3847/1538-3881/adb8e7.
- N. Jeffrey and B. D. Wandelt. *Machine Learning: Science and Technology*, 5(1):015008, Mar. 2024. doi: 10.1088/2632-2153/ad1a4d.
- W. D. Jennings, C. A. Watkinson, F. B. Abdalla, and J. D. McEwen. *MNRAS*, 483(3):2907–2922, Mar. 2019. doi: 10.1093/mnras/sty3168.
- D. Jimenez Rezende and S. Mohamed. *arXiv e-prints*, art. arXiv:1505.05770, May 2015. doi: 10.48550/arXiv.1505.05770.
- J. Kennedy et al. *Monthly Notices of the Royal Astronomical Society*, 529(4):3684–3698, Mar. 2024. ISSN 1365-2966. doi: 10.1093/mnras/stae760. URL <http://dx.doi.org/10.1093/mnras/stae760>.
- N. Kern. *MNRAS*, 541(2):687–713, Aug. 2025. doi: 10.1093/mnras/staf1007.
- N. S. Kern and A. Liu. *MNRAS*, 501(1):1463–1480, Feb. 2021. doi: 10.1093/mnras/staa3736.
- N. S. Kern et al. *ApJ*, 848(1):23, Oct. 2017. doi: 10.3847/1538-4357/aa8bb4.
- N. S. Kern et al. *ApJ*, 890(2):122, Feb. 2020. doi: 10.3847/1538-4357/ab67bc.
- J. R. Kerrigan et al. *ApJ*, 864(2):131, Sept. 2018. doi: 10.3847/1538-4357/aad8bb.
- D. P. Kingma and M. Welling. *arXiv e-prints*, art. arXiv:1312.6114, Dec. 2013. doi: 10.48550/arXiv.1312.6114.

- D. P. Kingma and M. Welling. Auto-encoding variational bayes, 2022. URL <https://arxiv.org/abs/1312.6114>.
- D. Korber, M. Bianco, E. Tolley, and J.-P. Kneib. *MNRAS*, 521(1):902–915, Feb. 2023. ISSN 1365-2966. doi: 10.1093/mnras/stad615. URL <http://dx.doi.org/10.1093/mnras/stad615>.
- Y. Kwon, S. E. Hong, and I. Park. *Journal of Korean Physical Society*, 77(1):49–59, July 2020. doi: 10.3938/jkps.77.49.
- P. La Plante and M. Ntampaka. *ApJ*, 880(2):110, Aug. 2019. doi: 10.3847/1538-4357/ab2983.
- O. Lahav. Deep machine learning in cosmology: Evolution or revolution?, 2023. URL <https://arxiv.org/abs/2302.04324>.
- Y. Le Cun, L. Bottou, and Y. Bengio. In *International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 151–154, Munich, 1997. IEEE. URL <http://leon.bottou.org/papers/lecun-97>.
- Y. LeCun. *Open Review*, 62(1):1–62, 2022.
- S. A. K. Leeney et al. *arXiv e-prints*, art. arXiv:2504.16791, Apr. 2025. doi: 10.48550/arXiv.2504.16791.
- P. Lemos, A. Coogan, Y. Hezaveh, and L. Perreault-Levasseur. *40th International Conference on Machine Learning*, 202:19256–19273, Jan. 2023. doi: 10.48550/arXiv.2302.03026.
- P. Lemos et al. *arXiv e-prints*, art. arXiv:2402.04355, Feb. 2024. doi: 10.48550/arXiv.2402.04355.
- W. Li et al. *Advances in Space Research*, 76(7):3768–3783, Oct. 2025. doi: 10.1016/j.asr.2025.05.011.
- X.-W. Liu, C. Heneka, and L. Amendola. *JCAP*, 2020(5):038, May 2020. doi: 10.1088/1475-7516/2020/05/038.
- J.-M. Lueckmann, G. Bassetto, T. Karaletsos, and J. H. Macke. *arXiv e-prints*, art. arXiv:1805.09294, May 2018. doi: 10.48550/arXiv.1805.09294.
- Y. Maan, J. van Leeuwen, and D. Vohl. *A&A*, 650:A80, June 2021. doi: 10.1051/0004-6361/202040164.
- Y. Mahida et al. *JCAP*, 2025(12):055, Dec. 2025. doi: 10.1088/1475-7516/2025/12/055.
- T. L. Makinen et al. *arXiv e-prints*, art. arXiv:2410.07548, Oct. 2024. doi: 10.48550/arXiv.2410.07548.
- T. Mangena, S. Hassan, and M. G. Santos. *MNRAS*, 494(1):600–606, May 2020. doi: 10.1093/mnras/staa750.
- A. Mazumder, A. Datta, A. Chakraborty, and S. Majumdar. *MNRAS*, 515(3):4020–4037, Sept. 2022. doi: 10.1093/mnras/stac1994.
- G. Mellema et al. *arXiv preprint arXiv:1501.04203*, 2015.

- R. Meriot and B. Semelin. *arXiv e-prints*, art. arXiv:2310.02684, Oct. 2023. doi: 10.48550/arXiv.2310.02684.
- R. Meriot, B. Semelin, and D. Cornu. *A&A*, 698:A80, June 2025. doi: 10.1051/0004-6361/202452901.
- F. G. Mertens, A. Ghosh, and L. V. E. Koopmans. *MNRAS*, 478(3):3640–3652, Aug. 2018. doi: 10.1093/mnras/sty1207.
- F. G. Mertens et al. *MNRAS*, 493(2):1662–1685, Apr. 2020. doi: 10.1093/mnras/staa327.
- F. G. Mertens, J. Bobin, and I. P. Carucci. *MNRAS*, 527(2):3517–3531, Jan. 2024. doi: 10.1093/mnras/stad3430.
- F. G. Mertens et al. *A&A*, 698:A186, June 2025. doi: 10.1051/0004-6361/202554158.
- M. Mesarcik, A.-J. Boonstra, E. Ranguelova, and R. V. van Nieuwpoort. *MNRAS*, 516(4):5367–5378, Nov. 2022. doi: 10.1093/mnras/stac2503.
- A. Mesinger and S. Furlanetto. *ApJ*, 669:663–675, Nov. 2007. doi: 10.1086/521806.
- A. Mesinger, S. Furlanetto, and R. Cen. *MNRAS*, 411:955–972, Feb. 2011. doi: 10.1111/j.1365-2966.2010.17731.x.
- R. Mondal, G. Mellema, S. G. Murray, and B. Greig. *MNRAS*, 514(1):L31–L35, July 2022. doi: 10.1093/mnras/rlac053.
- M. F. Morales et al. *MNRAS*, 483(2):2207–2216, Feb. 2019. doi: 10.1093/mnras/sty2844.
- O. Mosiane, N. Oozeer, and B. Bassett. pages 1–2, 10 2016. doi: 10.1109/RADIO.2016.7772036.
- B. P. Moster, T. Naab, M. Lindström, and J. A. O’Leary. *Monthly Notices of the Royal Astronomical Society*, 507(2):2115–2136, June 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab1449. URL <http://dx.doi.org/10.1093/mnras/stab1449>.
- S. Munshi et al. *A&A*, 681:A62, Jan. 2024. doi: 10.1051/0004-6361/202348329.
- S. Munshi et al. *arXiv e-prints*, art. arXiv:2507.10533, July 2025. doi: 10.48550/arXiv.2507.10533.
- S. Murray et al. *The Journal of Open Source Software*, 5(54):2582, Oct. 2020. doi: 10.21105/joss.02582.
- R. Neal. MCMC Using Hamiltonian Dynamics. In *Handbook of Markov Chain Monte Carlo*, pages 113–162. 2011. doi: 10.1201/b10905.
- S. Neutsch, C. Heneka, and M. Brüggén. *MNRAS*, 511(3):3446–3462, Apr. 2022. doi: 10.1093/mnras/stac218.
- S. Ni, Y. Li, L.-Y. Gao, and X. Zhang. *Astrophys. J.*, 934(1):83, 2022. doi: 10.3847/1538-4357/ac7a34.
- S. Ni et al. Pi-astrodeconv: A physics-informed unsupervised learning method for astronomical image deconvolution, 2024. URL <https://arxiv.org/abs/2403.01692>.

- S. Ni et al. *The Astrophysical Journal*, 990(2):122, 2025.
- G. M. Nita and D. E. Gary. *MNRAS*, 406(1):L60–L64, July 2010. doi: 10.1111/j.1745-3933.2010.00882.x.
- P. Ocvirk et al. *Monthly Notices of the Royal Astronomical Society*, 463(2):1462–1485, Aug. 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw2036. URL <http://dx.doi.org/10.1093/mnras/stw2036>.
- P. Ocvirk et al. *MNRAS*, 496(4):4087–4107, Aug. 2020. doi: 10.1093/mnras/staa1266.
- A. Offringa et al. *Monthly Notices of the Royal Astronomical Society*, 444(1):606–619, 2014.
- A. R. Offringa. AOFlagger: RFI Software. Astrophysics Source Code Library, record ascl:1010.017, Oct. 2010.
- A. R. Offringa et al. *MNRAS*, 405(1):155–167, June 2010. doi: 10.1111/j.1365-2966.2010.16471.x.
- A. Ore, C. Heneka, and T. Plehn. *SciPost Physics*, 18(5):155, May 2025. doi: 10.21468/SciPostPhys.18.5.155.
- M. Pagano et al. *MNRAS*, 520(4):5552–5572, Apr. 2023. doi: 10.1093/mnras/stad441.
- S. K. Pal et al. *Journal of Cosmology and Astroparticle Physics*, 2025(10):096, oct 2025. doi: 10.1088/1475-7516/2025/10/096. URL <https://doi.org/10.1088/1475-7516/2025/10/096>.
- S. K. Pal, A. Datta, and A. Mazumder. *JCAP*, 2025(2):058, Feb. 2025. doi: 10.1088/1475-7516/2025/02/058.
- G. Papamakarios, T. Pavlakou, and I. Murray. *arXiv e-prints*, art. arXiv:1705.07057, May 2017. doi: 10.48550/arXiv.1705.07057.
- J. Park, A. Mesinger, B. Greig, and N. Gillet. *MNRAS*, 484(1):933–949, Mar. 2019. doi: 10.1093/mnras/stz032.
- L. Parker et al. *MNRAS*, 531(4):4990–5011, July 2024. doi: 10.1093/mnras/stae1450.
- L. Parker et al. *arXiv*, art. arXiv:2510.17960, Oct. 2025a. doi: 10.48550/arXiv.2510.17960.
- L. Parker et al. *arXiv e-prints*, art. arXiv:2510.17960, Oct. 2025b.
- Y. Pietschke et al. *JCAP*, 2025(10):039, Oct. 2025. doi: 10.1088/1475-7516/2025/10/039.
- D. Piras et al. *The Open Journal of Astrophysics*, 7:73, Sept. 2024. doi: 10.33232/001c.123368.
- J. C. Pober et al. *ApJ*, 782(2):66, Feb. 2014. doi: 10.1088/0004-637X/782/2/66.
- S. Pochinda, J. Dhandha, A. Fialkov, and E. de Lera Acedo. Cosmological super-resolution of the 21-cm signal, 2025. URL <https://arxiv.org/abs/2502.00852>.
- P. R. Posture, Y. Mahida, S. Majumdar, and L. Noble. *arXiv e-prints*, art. arXiv:2510.01121, Oct. 2025.

- D. Prelogović and A. Mesinger. *MNRAS*, 524(3):4239–4255, Sept. 2023a. doi: 10.1093/mnras/stad2027.
- D. Prelogović and A. Mesinger. *MNRAS*, 524(3):4239–4255, Sept. 2023b. doi: 10.1093/mnras/stad2027.
- D. Prelogović and A. Mesinger. *A&A*, 688:A199, Aug. 2024. doi: 10.1051/0004-6361/202449309.
- D. Prelogović et al. *MNRAS*, 509(3):3852–3867, Jan. 2022. doi: 10.1093/mnras/stab3215.
- U. Rau and T. J. Cornwell. *Astronomy & Astrophysics*, 532:A71, 2011.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. URL <https://arxiv.org/abs/1505.04597>.
- N. Sabti et al. A generative modeling approach to reconstructing 21-cm tomographic data, 2024. URL <https://arxiv.org/abs/2407.21097>.
- M. G. Santos et al. *MNRAS*, 406(4):2421–2432, Aug. 2010. doi: 10.1111/j.1365-2966.2010.16898.x.
- A. Saxena et al. *MNRAS*, 525(4):6097–6111, Nov. 2023. doi: 10.1093/mnras/stad2659.
- A. Saxena et al. *RAS Techniques and Instruments*, 3(1):724–736, Jan. 2024. doi: 10.1093/rasti/rzae047.
- D. Schaurecker et al. Super-resolving dark matter halos using generative deep learning, 2022. URL <https://arxiv.org/abs/2111.06393>.
- K. Schmidt et al. *A&A*, 664:A134, Aug. 2022. ISSN 1432-0746. doi: 10.1051/0004-6361/202142113. URL <http://dx.doi.org/10.1051/0004-6361/202142113>.
- C. J. Schmit and J. R. Pritchard. *MNRAS*, 475(1):1213–1223, Mar. 2018. doi: 10.1093/mnras/stx3292.
- B. Schosser, C. Heneka, and T. Plehn. *SciPost Physics Core*, 8(2):037, Apr. 2025a. doi: 10.21468/SciPostPhysCore.8.2.037.
- B. Schosser, C. Heneka, and B. M. Schäfer. *arXiv e-prints*, art. arXiv:2508.10094, Aug. 2025b. doi: 10.48550/arXiv.2508.10094.
- A. K. Shaw et al. *JCAP*, 2021(12):024, Dec. 2021. doi: 10.1088/1475-7516/2021/12/024.
- B. C. Sherwin, S. S. Sethuram, C. Brummel-Smith, and J. H. Wise. *Research Notes of the AAS*, 7(11):242, nov 2023. doi: 10.3847/2515-5172/ad0cc5. URL <https://dx.doi.org/10.3847/2515-5172/ad0cc5>.
- H. Shimabukuro and B. Semelin. *MNRAS*, 468(4):3869–3877, July 2017. doi: 10.1093/mnras/stx734.
- P. H. Sims et al. *arXiv e-prints*, art. arXiv:2504.09725, Apr. 2025. doi: 10.48550/arXiv.2504.09725.
- SKA foregrounds chapter. *arXiv preprint arXiv:NNNN.MMMMM*, 2025.

SKA imaging chapter. *arXiv preprint arXiv:NNNN.MMMMM*, 2025.

O. M. Smirnov. *A&A*, 527:A106, Mar. 2011. doi: 10.1051/0004-6361/201016082.

J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.

K. Sooknunan et al. *arXiv e-prints*, art. arXiv:2412.15893, Dec. 2024. doi: 10.48550/arXiv.2412.15893.

A. Spurio Mancini et al. *Monthly Notices of the Royal Astronomical Society*, 511(2):1771–1788, Jan 2022. ISSN 1365-2966. doi: 10.1093/mnras/stac064. URL <http://dx.doi.org/10.1093/mnras/stac064>.

D. Sullivan, I. T. Iliev, and K. L. Dixon. *Monthly Notices of the Royal Astronomical Society*, 473(1):38–58, 09 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx2324. URL <https://doi.org/10.1093/mnras/stx2324>.

S. Talts et al. *arXiv e-prints*, art. arXiv:1804.06788, Apr. 2018. doi: 10.48550/arXiv.1804.06788.

E. Th  lie, F. Del Balso, J. B. Mu  oz, and A. Liu. *Phys. Rev. D*, 111(12):123501, June 2025. doi: 10.1103/jfkk-7q7l.

H. Tiwari et al. *JCAP*, 2022(4):045, Apr. 2022. doi: 10.1088/1475-7516/2022/04/045.

M. Trebitsch et al. *Monthly Notices of the Royal Astronomical Society*, 470(1):224–239, May 2017. ISSN 1365-2966. doi: 10.1093/mnras/stx1060. URL <http://dx.doi.org/10.1093/mnras/stx1060>.

M. Trebitsch et al. *A&A*, 653:A154, Sept. 2021. ISSN 1432-0746. doi: 10.1051/0004-6361/202037698. URL <http://dx.doi.org/10.1051/0004-6361/202037698>.

A. Tripathi, A. Datta, M. Choudhury, and S. Majumdar. *Monthly Notices of the Royal Astronomical Society*, 528(2):1945–1964, 01 2024. ISSN 0035-8711. doi: 10.1093/mnras/stae078. URL <https://doi.org/10.1093/mnras/stae078>.

A. Tripathi, G. Kaur, A. Datta, and S. Majumdar. *JCAP*, 2024(10):041, Oct. 2024. doi: 10.1088/1475-7516/2024/10/041.

A. Tripathi, A. Datta, A. Mazumder, and S. Majumdar. *JCAP*, 2025(10):035, Oct. 2025. doi: 10.1088/1475-7516/2025/10/035.

A. Vafaei Sadr et al. *MNRAS*, 499(1):379–390, Nov. 2020. doi: 10.1093/mnras/staa2724.

F. Villaescusa-Navarro et al. *ApJ*, 929(2):132, Apr. 2022. doi: 10.3847/1538-4357/ac5d3f.

F. Villaescusa-Navarro et al. *arXiv*, art. arXiv:2510.26887, Oct. 2025. doi: 10.48550/arXiv.2510.26887.

H. Wang et al. *Phys. Rev. D*, 111(10):103531, 2025. doi: 10.1103/PhysRevD.111.103531.

- S. Wyithe, P. M. Geil, and H. Kim. *PoS*, AASKA14:015, 2015. doi: 10.22323/1.215.0015.
- Z. Yang, C. Yu, J. Xiao, and B. Zhang. *MNRAS*, 492(1):1421–1431, Feb. 2020. doi: 10.1093/mnras/stz3521.
- S. Yatawatta, L. De Clercq, H. Spreeuw, and F. Diblen. *arXiv e-prints*, art. arXiv:1904.05619, Apr. 2019. doi: 10.48550/arXiv.1904.05619.
- Y.-J. Zhang, Y.-Z. Li, J. Cheng, and Y.-H. Yan. *Research in Astronomy and Astrophysics*, 21(12):299, Dec. 2021. doi: 10.1088/1674-4527/ac2944.
- X. Zhao, Y. Mao, C. Cheng, and B. D. Wandelt. *ApJ*, 926(2):151, Feb. 2022a. doi: 10.3847/1538-4357/ac457d.
- X. Zhao, Y. Mao, and B. D. Wandelt. *ApJ*, 933(2):236, July 2022b. doi: 10.3847/1538-4357/ac778e.
- X. Zhao, Y.-S. Ting, K. Diao, and Y. Mao. *MNRAS*, 526(2):1699–1712, Dec. 2023. doi: 10.1093/mnras/stad2778.
- X. Zhao, Y. Mao, S. Zuo, and B. D. Wandelt. *ApJ*, 973(1):41, Sept. 2024. doi: 10.3847/1538-4357/ad5ff0.
- X. Zhao et al. *ApJ*, 998(2):189, Feb. 2026. doi: 10.3847/1538-4357/ae2c7e.
- Y. Zhou and P. La Plante. *PASP*, 134(1034):044001, Apr. 2022. doi: 10.1088/1538-3873/ac5f5d.