



Methodological Frontiers in 21-cm Intensity Mapping: the Treatment of Systematics and Foreground Contamination

Marta Spinelli^{1,2} Matilde Barberi-Squarotti^{3,8,10,23} José Luis Bernal⁴ Somnath Bharadwaj⁵ Victor Bonjean⁶ Philip Bull^{7,2} Carmelita Carbone⁸ Isabella P. Carucci⁹ Stefano Camera¹⁰ Suman Chatterjee² Tianyue Chen⁷ Zhaoting Chen¹¹ Samir Choudhuri¹² Devin Crichton¹³ Steven Cunnington^{14,7} Bianca De Caro⁸ Khandakar Md Asif Elahi¹² Jose Fonseca^{15,22,2} Athanasia Gkogkou⁶ Abhik Ghosh¹⁶ Keith Grainge⁷ Wenkai Hu² Melis O. Irfan^{17,2} Yichao Li¹⁸ Siyambonga Matshawule² Geoff Murphy² Simon Prunet¹ Mario G. Santos^{2,19} Jean-Luc Starck²⁰ Jingying Wang^{21,2} Laura Wolz⁷ and Zheng Zhang⁷

¹*Observatoire de la Côte d'Azur, Laboratoire Lagrange, Bd de l'Observatoire, 06304 Nice, France*

²*Department of Physics & Astronomy, University of the Western Cape, Cape Town 7535, South Africa*

³*Dipartimento di Fisica, Università degli Studi di Milano, via G. Celoria 16, 20133 Milano, Italy*

⁴*Instituto de Física de Cantabria (IFCA), CSIC-Univ. de Cantabria, Avda. de los Castros s/n, E-39005 Santander, Spain*

⁵*Department of Physics, Indian Institute of Technology Kharagpur, Kharagpur 721 302, India*

⁶*Institutes of Computer Science and Astrophysics, Foundation for Research and Technology Hellas (FORTH), Greece*

⁷*Jodrell Bank Centre for Astrophysics, Department of Physics & Astronomy, The University of Manchester, Manchester M13 9PL, UK*

⁸*INAF, Osservatorio Astrofisico di Brera-Merate, via Brera 28, 20121 Milano, Italy*

⁹*INAF, Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, 34131 Trieste, Italy*

¹⁰*Dipartimento di Fisica, Università degli Studi di Torino, via P. Giuria 1, 10125 Torino, Italy*

¹¹*Institute for Astronomy, The University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK*

¹²*Centre for Strings, Gravitation and Cosmology, Department of Physics, Indian Institute of Technology Madras, Chennai 600036, India*

¹³*Institute of Particle Physics & Astrophysics, Department of Physics, ETH Zurich, 8093 Zurich, Switzerland*

¹⁴*Institute of Cosmology & Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, UK*

¹⁵*Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto CAUP, Rua das Estrelas, 4150-762 Porto, Portugal*

¹⁶*Department of Physics, Banwarilal Bhalotia College, Asansol, West Bengal-713303, India*

¹⁷*Institute of Astrophysics, University of Cambridge, Madingley Road, CB3 0HA*

¹⁸*Department of Physics, College of Sciences, Northeastern University, Wenhua Road, Shenyang, China*

¹⁹*South African Radio Astronomy Observatory (SARAO), Cape Town, 7700, South Africa*

²⁰*Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191, Gif-sur-Yvette, France*

²¹*Shanghai Astronomical Observatory, Chinese Academy of Sciences, 80 Nandan Road, Shanghai 200030, China*

²²*Departamento de Física e Astronomia, Faculdade de Ciências, Universidade do Porto, Rua do Campo Alegre 687, 4169-007 Porto, Portugal*

²³*INFN – Istituto Nazionale di Fisica Nucleare, Sezione di Milano, via G. Celoria 16, 20133 Milano, Italy*

E-mail: marta.spinelli@oca.eu

The distribution of neutral hydrogen (H_I) in the post-reionization universe traces the cosmic large-scale structure and therefore serves as a powerful cosmological probe. An efficient way to measure its distribution over large sky areas and wide redshift ranges is through single-dish intensity mapping, which exploits the autocorrelation signal of each dish in a telescope array, while they are scanning the same patch in the sky. Thanks to its broad frequency coverage and technical capabilities, SKA-Mid will enable measurements of the integrated 21 cm emission from H_I up to redshift $z \sim 3$, making single-dish intensity mapping a key observable for probing dark matter and dark energy.

Isolating the faint 21 cm cosmological signal without introducing distortions or biases is, however, a challenging task. The 21 cm signal is several orders of magnitude weaker than the astrophysical foregrounds and its analysis is further affected by instrumental systematics. Overcoming these difficulties requires detailed modelling, as well as continuous improvements and innovations in data analysis techniques. Over the past decade, the international community has been developing and testing new methods, both to address current observational challenges and to prepare for the forthcoming SKA-Mid observations.

This chapter presents a review of recent advances in map-making and component separation techniques, with particular emphasis on dealing with telescope-specific systematics such as beam response and correlated noise. We focus on results obtained in controlled simulation environments, providing a valuable framework for assessing the strengths and limitations of different approaches.

Developing robust algorithms capable of accurately handling instrumental effects and sky-model uncertainties is a crucial step toward fully exploiting the cosmological potential of H_I intensity mapping surveys in the SKA Observatory era.

1 Introduction

Intensity Mapping (IM) is a technique to trace the cosmic distribution of neutral hydrogen directly from the integrated luminosity of its redshifted 21 cm line, without the need for the detection of individual sources. Despite the sensitivity performance of the SKA Observatory (SKAO) (Braun et al., 2019), the detection of neutral hydrogen clouds beyond the local Universe will remain a challenge, and only a small number of very bright galaxies at 21 cm will be detected at $z \sim 1$. IM will instead significantly extend the redshift range accessible for observing the large-scale distribution of neutral hydrogen, thus accessing the underlying matter distribution as a function of cosmic time (e.g. Wyithe and Loeb, 2008; Bagla et al., 2010; Ansari et al., 2012; Battye et al., 2013).

With the SKA-Mid precursor MeerKAT, the MeerKLASS collaboration has pioneered the approach of performing IM operating the entire array as a fast-scanning collection of multiple individual dishes and exploiting the autocorrelation signals, also known as “single-dish” mode (Santos et al., 2017). This enabled the mapping of the sky on angular scales that the interferometer would resolve out, while using one of MeerKAT’s strengths: the speed and precision of the antennas’ synchronized movement. A summary of the work and results achieved via single-dish IM on MeerKAT in recent years is presented in chapter Cunnington&Wang et al. (2026). These results show that the technique, originally proposed as an idea in the first edition of this book about 10 years ago (Santos et al., 2015), has evolved into a concrete reality. The experiences gained with MeerKAT, including lessons learned, challenges overcome, and best practices developed, provide a foundation for SKA-Mid to achieve even more robust results, marking a significant step toward the full maturity of radio cosmology.

Indeed, a single-dish survey with SKA-Mid for the post-reionization universe will allow us to place tight constraints on cosmological parameters that are both competitive and complementary to current galaxy surveys, challenging our understanding of dark matter and dark energy. A detailed description of the anticipated impact on cosmology of an intensity mapping survey with SKA-Mid is described in detail in Wolz et al. (2026); Camera et al. (2026); Fonseca et al. (2026), where the interested reader can also find a carefully compiled list of references. A summary of the methodologies developed to capture and exploit non-Gaussian features in 21 cm signals from future SKA-Mid observations via various higher-order statistics can instead be found in Majumdar et al. (2026).

The process of transforming the autocorrelation signals from each SKA-Mid antenna observing the same patch of sky, into a single 21 cm data cube that accurately represents the cosmic 21 cm signal and is ready for cosmological interpretation, will not be trivial. It will require constant attention to the response and instrumental characteristics of the telescope, as well as continuous improvements and innovations in analysis techniques to handle contaminants and systematic effects. The interested reader can find a detailed description of the Radio Frequency Interference (RFI) flagging, calibration, and map-making pipeline currently developed for MeerKAT in Wang et al. (2021) and MeerKLASS Collaboration et al. (2025).

Another critical challenge, coupled with calibration and map-making, lies in isolating the faint 21

cm signal without distorting or biasing its intrinsic cosmological information. The cosmological H_I emission is buried beneath much brighter foreground radio emissions that dominate the same frequency range. The strongest of these, synchrotron emission from our own Galaxy, exceeds the expected 21 cm signal by roughly five orders of magnitude, even at high Galactic latitudes. Similarly, extragalactic point sources can be about three orders of magnitude brighter. Such strong foregrounds can severely contaminate the data, unless they are removed with exceptional precision.

In the last decade, many studies have addressed the problem in the context of single-dish H_I IM and investigated the quality of foreground removal methods on the data (Switzer et al., 2015; Wolz et al., 2017; Cunnington et al., 2020; Carucci et al., 2025), as well as simulations (e.g., Ansari et al., 2012; Wolz et al., 2014; Alonso et al., 2014; Shaw et al., 2015; Olivari et al., 2016; Carucci et al., 2020; Mäkinen et al., 2021; Yohana et al., 2021; Fonseca and Liguori, 2021; Soares et al., 2021; Spinelli et al., 2022), where blind and non-parametric methods have proven to be the most powerful.

A further complication arises from the instrument itself, which can introduce additional systematics such as the telescope beam response, amplitude gain fluctuations, polarization leakage, internal signal-chain reflections, and digital non-linearities. Overcoming or at least mitigating these effects requires meticulous calibration and a deep understanding of instrumental behaviour.

The telescope's primary beam, for example, affects both calibration and component separation, imprinting its own frequency dependence on the foregrounds. Accurately characterizing the beam, particularly far from its centre, is difficult due to variations between individual dishes and temporal changes. Understanding the extent to which these instrumental effects may hinder the detection of the cosmological signal is therefore of paramount importance.

The controlled framework of simulations allows the investigation of these individual observational challenges of H_I IM. Studies related to SKAO, its precursors, and pathfinders include the primary beam effects (Matshawule et al., 2021; Spinelli et al., 2022; Gkogkou et al., 2026), polarization leakage (Shaw et al., 2015; Carucci et al., 2020; Cunnington et al., 2021), $1/f$ noise (Harper et al., 2018; Chen et al., 2020; Li et al., 2021; Irfan et al., 2024) and radio frequency interference due to satellites (Harper et al., 2018; Engelbrecht et al., 2025).

The chapter is organized as follows. In section 2, we provide a very brief overview of the typical choices for the simulated single-dish IM observations. We present a brief review of the problem of component separation in the context of intensity mapping highlighting the difference among blind and non-blind methods in section 3. In section 4, we focus on blind component separation techniques providing a brief review of the various methodologies discussed in literature and typically developed for solving a similar problem in the context of the Cosmic Microwave Background (CMB). We summarize the finding of a component separation data challenge that we performed within the Cosmology Science Working Group (SWG) in section 5. In section 6, we deal more specifically with the added complexity of a primary beam sidelobes and frequency evolution. In section 7 we explore instead component separation methods that relies on strategies different from blind methods, including the the multi-frequency angular power spectrum approach, AI driven solutions and treating the problem in higher dimensional space using Bayesian statistics. In section 8, we discuss the issue of signal loss in component separation. In section 10 we focus on the $1/f$, on how

this can be characterized and dealt with at map-making process. We conclude in [section 11](#) where we provide a brief final review on other possible types of systematics and the challenges that the community will need to tackle for to fully exploit the capabilities of SKA-Mid for precision large scale structure studies with H_I IM.

2 Typical choices for IM simulations

A typical IM mock data cube contains: i) foreground emission that are bright in the frequency range of interest (SKA-Mid Band 1 and 2), ii) simulated temperature brightness from the 21 cm line at the corresponding redshifts (up to $z \sim 3$ for Band 1), iii) a simplified description of the effect of the telescope beam, and iv) a white noise realisation from the radiometer equation. Since the noise level decreases with the square root of the number of dishes, the SKA-Mid AA4 configuration, with 197 dishes, is expected to achieve a lower noise level than the AA* configuration, which includes 144 dishes. Therefore, AA4 represents our ideal-case scenario.

The astrophysical foregrounds often included in the mock data are galactic synchrotron and free-free emissions, and extragalactic radio point sources. The synchrotron and free-free emissions are generally modelled as smooth power laws in frequency, with amplitudes and spectral indices provided by the Planck Legacy Archive FFP10 simulations or the Planck Sky Model ([Delabrouille et al., 2013](#)) (in particular via the popular *pysm* software ([Thorne et al., 2017](#))). For extragalactic radio point sources (i.e., active galactic nuclei, radio galaxies, quasars, etc.), empirical models (e.g. [Battye et al., 2013](#)) or information from available surveys (e.g. [Matshawule et al., 2021](#)) are used. Simulated leakage of polarized foreground to total intensity has been also considered often using the models from [Alonso et al. \(2014\)](#); [Shaw et al. \(2015\)](#). The polarised signal could have frequency structure due to Faraday rotation ([Spinelli et al., 2018](#)) complicating the IM analysis.

The cosmological H_I maps are normally computed using lognormal realisation of the matter field and then transformed into temperature brightness assuming a scale-independent H_I bias. A popular software is described in [Alonso et al. \(2014\)](#). With the progress of the data quality and the data analysis tools in the last decade, the need for more realistic simulations that could help the interpretation of the signal has become more pressing. A review of some of the available simulations that are used by the IM community can be found in [Ronconi et al. \(2026\)](#). Simulations including more sophisticated H_I treatment and galaxy properties are useful in particular for cross-correlation analysis with galaxy surveys (e.g. [Wolz et al., 2021](#)).

To mimic an SKA-Mid observation, the maps are often smoothed with a frequency-dependent Gaussian telescope beam (which we will discuss in more detail in [section 6](#)) and add Gaussian thermal noise, following the prescriptions of [SKA Cosmology SWG et al. \(2020\)](#).

3 Foreground cleaning methodology

There are a variety of strategies to deal with these extragalactic and Galactic foreground contaminations. Broadly speaking, one can categorize them into a spectrum of methods going from foreground avoidance, to blind foreground subtraction, to semi-blind or non-blind methods, where some level of knowledge on the components is imposed in the solution.

Foreground avoidance finds an observational window at the power spectrum level in $k_{||} - k_{\perp}$ space, where foregrounds do not substantially affect the estimates of the power spectrum (Shaw et al., 2015, e.g.). *Blind foreground subtraction* methods, that are described in more detail in section 4, have been the most successful in the last years, especially facing real data. There are also *methods involving some knowledge of the foregrounds*. Some earlier attempts to subtract foregrounds imposed a polynomial logarithmic structure in frequency (Bowman et al., 2009). Others have used MeerKLASS data to calibrate the spectral index of galactic synchrotron (Irfan et al., 2022) and gain insight of the foregrounds which can later be used on foreground removal. In Fonseca and Liguori (2021), they used the models of Santos et al. (2005) to study if one could marginalize over the foregrounds models and learn about their angular structure. These parametric approaches are not the sole methods using prior knowledge of contaminants: the authors of Murphy et al. (2024) model some of the instrumental effects which are then marginalized over using an Hamiltonian Monte Carlo sampler.

4 Blind component separation

Several widely used and well-tested methods for addressing the challenge of separating contaminants in 21 cm IM rely only on two main assumptions: i) foreground emissions vary slowly with frequency compared to the rapidly oscillating cosmological signal and, ii) foreground emission is significantly larger in magnitude than the H_I signal. We can identify the primary foreground contributions as a set of smooth functions of frequency, which can be subtracted from the data to isolate the H_I signal.

In practice, IM observations can be represented by two-dimensional (frequency and pixel) data-cubes, $\mathbf{X} \equiv T_{\text{obs}}(\nu, \hat{\mathbf{n}})$. It is then assumed that the matrix \mathbf{X} , dominated by the foreground emission, can be linearly decomposed into a set of N_{fg} “foreground” sources in pixel space \mathbf{S} modulated in frequency through a mixing matrix \mathbf{A} plus some residuals \mathbf{N} that should in principle contain most of the cosmological signal that we aim to recover plus the white instrumental noise:

$$\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{N}. \quad (1)$$

The process of contaminant separation involves solving Equation 1 to determine \mathbf{N} . The assumptions made to derive matrix \mathbf{A} and components \mathbf{S} that satisfy Equation 1 differ among various methods. In the next sections, we outline the main ideas following Spinelli et al. (2022). We refer interested readers to the various references discussed below for a more accurate treatment of the problem.

4.1 Principal Component Analysis

Principal Component Analysis (PCA) can be used to identify an estimate for the mixing matrix \mathbf{A} , the columns of which will be given by the first N_{fg} principal components. The principal components are essentially the eigenvectors of the mean-centred data $\nu\nu'$ covariance matrix \mathbf{C} , given by

$$C_{ij} = \frac{1}{N_{\hat{n}}} \sum_{p=1}^{N_{\hat{n}}} w_i \Delta T(\nu_i, \hat{n}_p) w_j \Delta T(\nu_j, \hat{n}_p), \quad (2)$$

where $\Delta T(v_i, \hat{n}_p) = T(v_i, \hat{n}_p) - \bar{T}(v_i)$ and the summation is over all $N_{\hat{n}}$ pixels. The w factors provide an optional map weighting. The eigendecomposition is given by $\mathbf{C}\mathbf{V} = \mathbf{V}\mathbf{\Lambda}$, where $\mathbf{\Lambda}$ is the diagonal matrix of N_v eigenvalues. The first N_{fg} columns from the eigenvector matrix \mathbf{V} represent the entries for the mixing matrix.

4.1.1 Multiscale PCA

It has recently been proposed to take advantage of PCA in a multiscale framework to complement with angular information its frequency-direction-only assumptions on the contaminants (Carucci et al., 2025). Each frequency two-dimensional map is first decomposed into wavelet-filtered large- and small-scale maps, using the isotropic undecimated wavelet transform (Starck et al., 2007). The original data cube now corresponds to two cubes, referring either to the large- or small-scale only information, which are then cleaned separately. The component separation process becomes two independent processes, with their own, independently defined mixing matrix and number of removed components, N_{fg} , resulting in a better characterization of the contaminants. The final cleaned map is assembled by summing the two independently cleaned small- and large-scale maps, since the wavelet decomposition is numerically exact.

4.2 Independent Component Analysis

Independent Component Analysis (ICA) is a blind source separation method that extracts statistically independent signals without prior knowledge of how they were combined. It assumes that the data are linear combinations of independent sources and identify them by maximizing any statistical quantity that measures non-Gaussianity. In particular, Fast Independent Component Analysis (FastICA), developed in Hyvärinen (1999), has been employed for foreground cleaning on simulated H_I data (Chapman et al., 2012; Wolz et al., 2014; Alonso et al., 2014; Cunnington et al., 2019; Carucci et al., 2020; Cunnington et al., 2021) as well as real data (Wolz et al., 2017, 2021; Hothi et al., 2021). Since it often contains a “whitening” step that can be achieved with a PCA analysis, FastICA is essentially an extension of PCA, and hence in most cases in the context of foreground cleaning, will provide very similar results.

4.3 Generalised Morphological Component Analysis

Generalised Morphological Component Analysis (GMCA) is a blind component separation method based on sparsity (Bobin et al., 2007) that has been successfully applied in various astrophysical contexts. It assumes that the N_{fg} foreground components verify two hypotheses: they are sparse in a given transformed domain (i.e., most samples are zero-valued) and their supports are disjoint; in other words, the foreground components are *morphologically* diverse (i.e., their non-zero samples appear at different locations).

Carucci et al. (2020) showed the wavelet domain to be optimal to sparsely describe foregrounds and contaminants in the low- z H_I IM context. Firstly, we project the data \mathbf{X} onto wavelet space. The GMCA algorithm aims at minimizing the following cost function:

$$\min_{\mathbf{A}, \mathbf{S}} \sum_{i=1}^{N_{fg}} \lambda_i \|S_i\|_1 + \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2, \quad (3)$$

where the first term is the ℓ_1 norm, i.e. $\sum_{j,k} |S_{j,k}|$: this constitutes a constraint for sparsity, mediated by the regularization coefficients λ_i . The second term is the usual data-fidelity ℓ_2 norm term. We find solutions for \mathbf{A} and \mathbf{S} by iterating a projected alternate least-square procedure: we fix \mathbf{A} and perform a least-square update to determine \mathbf{S} , we compute the thresholds λ_i via mean absolute deviation of S_i , we update \mathbf{A} with \mathbf{S} fixed and so on. The key point is the thresholding: it makes possible to keep the samples with the highest amplitudes, which are the most informative to retrieve the mixing matrix \mathbf{A} (i.e., they most likely to belong to the foreground components and are the least likely to be contaminated by the cosmological signal and noise), and it provides robustness in terms of convergence since the thresholds decrease with the progressive iterations.

4.4 Generalized Needlet Internal Linear Combination

The Internal Linear Combination (ILC) technique was initially introduced in the analysis of WMAP CMB data (Bennett et al., 2003). It extracts one or more components with known spectral behavior from multi-frequency observations by applying a weight vector W . These weights are designed to preserve unit response to the target signal while minimizing the overall variance contributed by residual foregrounds.

Looking back at Equation 1, the ILC estimate of \mathbf{S} is obtained as a weighted linear combination, $\hat{\mathbf{S}} = \mathbf{W}^T \mathbf{X}$. A major advantage of the ILC approach is that it does not rely on explicit modeling of the foregrounds. Its main limitation, however, is that it can only recover components with a spatially invariant frequency dependence.

The Generalized Needlet ILC (GNILC) method extends the ILC framework to recover emissions that cannot be represented by a single template with fixed frequency scaling, such as diffuse foregrounds. It operates in needlet space and incorporates *prior* information on the power spectrum of the target signal to compensate for the lack of precise knowledge of its spectral dependence (Remazeilles et al., 2011). Needlets are a kind of spherical wavelets introduced for statistical analysis by Narcowich et al. (2006); their properties are discussed in Baldi et al. (2009) and Marinucci and Peccati (2011). Their first application in cosmology was for CMB data analysis (Pietrobon et al., 2006; Marinucci et al., 2008). The needlet decomposition provides simultaneous localization in both sky position and angular scale, enabling a more flexible component separation.

GNILC has been successfully applied to CMB analysis by Planck Collaboration et al. (2016), and subsequently adapted for 21 cm intensity mapping studies by Olivari et al. (2016) and the BINGO collaboration (Fornazier et al., 2022).

The GNILC algorithm requires two main inputs: the theoretical 21 cm power spectrum across frequency channels (in this sense it can be considered a *semi-blind* method) and an estimate of the number of foreground components, N_{fg} , which is needed to run a constrained PCA step to separate the HI and the foreground plus noise components. By transforming the covariance matrix of the observations using the information from the theoretical prior on the HI signal, the constrained PCA identifies the HI subspace as the eigenvectors whose eigenvalues are approximately equal to 1, while the remaining N_{fg} with larger eigenvalues define the foreground plus noise subspace.

This procedure is carried out independently at each needlet scale using the Akaike Information

Criterion (AIC) (Akaike, 1974). For each sky location and needlet scale, N_{fg} is determined by minimizing the AIC:

$$\text{AIC}(N_{\text{fg}}) = 2n N_{\text{fg}} - 2 \ln \mathcal{L}_{\text{max}}(N_{\text{fg}}), \quad (4)$$

where n is the number of modes within the needlet domain, and \mathcal{L}_{max} denotes the maximum-likelihood estimate of the data covariance matrix under a model with N_{fg} independent foreground components (see Appendix A of Olivari et al. (2016) for details). Once N_{fg} has been determined, the ILC weight matrix \mathbf{W} is computed at each needlet scale, allowing reconstruction of the H_I signal from the observational data.

In De Caro et al. (2025), the GNILC method was tested on simulated H_I intensity maps for SKAO-Mid in the AA4 configuration. The assumptions for such simulations are taken from Carucci et al. (2020), and are typical for most of the works discussed in this review chapter (see section 2). The telescope beam is assumed Gaussian, with $\theta_{\text{FWHM}} = \frac{c}{\nu D}$, where D is the dish diameter, and it varies across the frequency channels.

They also implement a needlet version of PCA (Need-PCA) and GMCA (Need-GMCA). Each PCA or GMCA step is applied to the needlet coefficients of the observation map after needlet decomposition. The number of foreground components to remove is set at the same N_{fg} , for each needlet coefficient. In the left panel of Figure 1, GNILC performance is compared with PCA and GMCA, both in their standard version and when using needlets, plotting the relative difference between the input H_I angular power spectrum and the recovered one. At larger scales ($\ell \lesssim 30$), GNILC underestimates the input power spectrum. At those scales, the cosmological signal and the foregrounds are degenerate; hence GNILC tends to remove more components, leading to smaller foreground residuals in the H_I plus noise maps compared to the other methods. Conversely, a fraction of the H_I plus noise is also removed along with the foregrounds, causing signal loss.

Projecting the mixing matrix onto the pure-foreground (e.g. Cunnington et al., 2021) one can estimate the foreground that leaks into the cosmological signal plus noise, as shown in the right panel of Figure 1. For the very large scales ($\ell < 15$), GNILC achieves lower foreground residuals than the other methods, as it dynamically adapts the number of foreground components depending on scale and sky region, whereas this is fixed in the other methods. As the galactic plane is excluded in the analysis, considering the polarization leakage following the procedure of Alonso et al. (2014) does not alter the conclusion. The same is true considering or not a frequency evolving beam (the standard λ/D dependency).

5 The foreground cleaning data challenge of the Intensity Mapping Focus Group

In Spinelli et al. (2022), a detailed study of foreground removal methods is presented in particular for AA4 configuration SKA-Mid H_I IM surveys. The work has been an initiative of the H_I Intensity Mapping Focus Group of the SKA Cosmology Science Working Group: a data challenge for component separation where participants were presented with simulated data-cubes of unknown foregrounds, H_I signal and instrumental specifics such as the beam and noise level.

Data challenges are a useful test for the maturity of analysis pipelines and have been routinely exploited by the SKA Observatory in recent years (Bonaldi and Braun, 2018; Bonaldi et al., 2020).

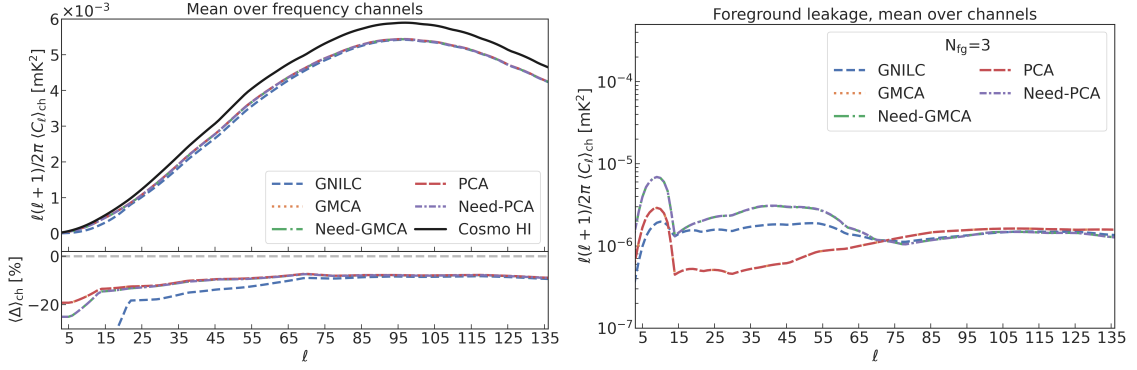


Figure 1: *Left.* Comparison between the C_ℓ^{Res} recovered with GNILC (blue line), GMCA (orange line), Need-GMCA (green line), PCA (red line) and Need-PCA (purple line). The input C_ℓ^{Cosmo} is plotted in black. The bottom panel shows the percent difference Δ . The telescope beam is Gaussian, with $\theta_{\text{FWHM}} = \frac{c}{\nu D}$, where D is the dish diameter, and it varies across the frequency channels. $N_{\text{fg}}=3$ for PCA, Need-PCA, GMCA and Need-GMCA; GNILC computes N_{fg} for each scales and region of pixel according to the AIC. *Right.* comparison between the C_ℓ of the foreground leakage obtained by averaging over the frequency channels. Figures taken from De Caro et al. (2025).

The realistic simulations for the task were constructed with two different implementations of the astrophysical foregrounds based on Santos et al. (2005) and the Planck Sky Model (Planck Collaboration et al., 2016), with a realistic scanning strategy for a $\sim 5000 \text{ deg}^2$ survey resulting in anisotropic noise, as well as a more sophisticated beam model (i.e. gaussian tapered Airy disk model presented in Harper and Dickinson 2018, Airy in short) with chromatic side-lobes for the SKA-Mid and MeerKAT dishes. A more detailed discussion on the beam can be found in section 6.

The true level of H I signal, obtained from a halo catalogue painted with a halo-H I mass relation derived from a galaxy evolution semi-analytical model. For more details on H I IM simulation in the context of the SKAO see Ronconi et al. (2026). The participants used a total of nine different pipelines to clean the data-cubes, ranging from different blind (PCA, fastICA and GMCA) and non-blind (parametric fitting) source separation algorithms. The results were analysed not only in terms of their absolute performance but also in terms of relative performance in order to understand better their weaknesses and strengths.

The results pointed out that, even among similar methods, subtleties related to each specific implementation can lead to substantial differences in the cleaning performance, and that the choice N_{fg} is not easily deducible and objective without extra prior information on the signal. In Figure 2, we report the performances of the various methods using the line-of-sight (los) power spectrum, a powerful statistic for residual structure in frequency. Blind methods are generally capable of recovering within 20% the input power spectra in the frequency range (left panel) and spatial scales with the least beam suppression, for a simple Gaussian beam. The right panel shows instead the case of a Airy beam whose frequency dependence is fitted on MeerKAT’s beam measurements. The presence of residual structure to to the coupling of the beam structure with spatial variation of the foreground was anticipated in Matshawule et al. (2021) and will be discuss in more detail in the next section.

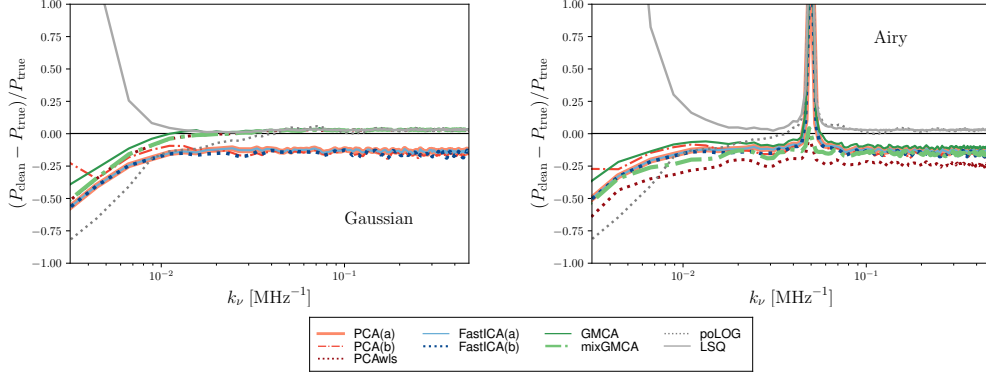


Figure 2: Comparison of the estimator $(P_{\text{clean}} - P_{\text{true}})/P_{\text{true}}$, where P_{clean} is the power spectrum of the residual maps for the various cleaning methods, while P_{true} is the input signal and noise. The left panel show the Gaussian beam case while the right panel show the result for a realistic beam model (Airy, see Harper and Dickinson 2018)

Aiming at a detailed comparison among methods, the performances in reconstructing the true signal have been summarized in metrics defined for both the los power spectrum and the angular power spectra at each frequency/redshift (see Spinelli et al. 2022 for details). These have been reported in the radar charts in Figure 3: a better performing method corresponds to a more extended chart. Methods are colour-coded as in Figure 2. For each blind method, the number of subtracted components N_{fg} is also reported, and the intensity of the colour is scaled proportionally (darker colour corresponds to higher N_{fg}). Note that mixGMCA is associated with two different N_{fg} the first for the largest scale PCA and the second for smaller scales GMCA. A similar multi-scale cleaning approach (see also subsection 4.1) has been recently successfully applied to real data in Carucci et al. (2025).

In the right panel the results are shown again for the realistic beam case but when the mock data are artificially resmoothed to a common low resolution via a Gaussian convolution. This is was a popular procedure for real data (e.g. Wolz et al., 2017) that was helping the cleaning procedure. We see here that the same is not true in simulations when there is a mismatch between the simulated telescope beam (Airy) and the one assumed for the resmoothing. A detailed discussion can be found also in Matshawule et al. (2021).

6 The effect of the telescope beam

In single-dish IM, the telescope beam is often assumed to be Gaussian with a FWHM, $\Delta\theta$, given approximately by

$$\Delta\theta \approx 1.16 \frac{\lambda}{D}, \quad (5)$$

where λ is the observed wavelength and D the dish diameter. In reality, there are two main complications to this scenario: i) the beam has sidelobes that can be more or less pronounced (see the left panel of Figure 4) and can introduce structures in the maps coming from strong point sources

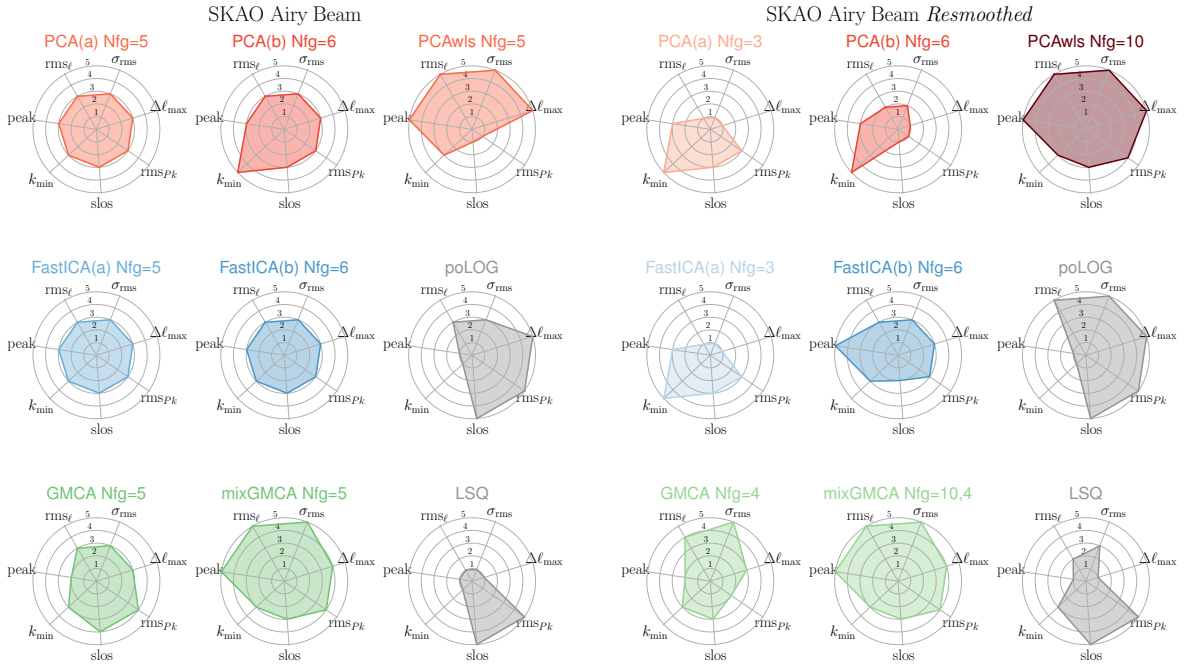


Figure 3: Radar charts showing the performance of the various cleaning methods participating in the data challenge presented in Spinelli et al. (2022) (larger covered areas correspond to better performing methods). Mock data are assuming a SKA-Mid H₁ IM survey. See text for details.

picked up by that; ii) the beam can have a non-trivial chromaticity (see for example the structure in the right panel of Figure 4) that can severely affect component separation.

6.1 Beam sidelobe structure

Sidelobes can be problematic for IM experiments. For example, the extragalactic point sources which are located in the sidelobes of the primary beam pattern, can show up as oscillations in the measured multifrequency angular power spectrum (MAPS) $C_\ell(\Delta\nu)$ (Ghosh et al., 2011a), which we will describe in more detail in subsection 7.1. These oscillations are extremely difficult to model and remove from the measured $C_\ell(\Delta\nu)$ but can be mitigated using window functions. This idea is implemented in the Tapered Gridded Estimator (TGE) again described in subsection 7.1. We refer the interested reader to Ghosh et al. (2011b); Choudhuri et al. (2014, 2016) for a detailed description.

In the specific context of blind component separation for single-dish IM, the results of Matshawule et al. (2021) seems to suggest that sidelobe structure of the primary beam do not cause major problems per se if there is no complex chromaticity in the beam. This can be seen in Figure 5 (black and yellow lines), where, until the variation of the FWHM is smooth with frequency, the level of the point source contamination does not change the performance of the foreground cleaning.




Figure 4: *Left:* Primary beam model at 950 MHz: Gauss model (black), Jinc model (Wilson and Rohlf s K., 2013) (solid green), the cosine model (Condon and Ransom, 2016) (magenta), the gaussian tapered Airy disk used in Harper and Dickinson (2018) (dashed red), and the one obtained from the Eidos package presented in Asad et al. (2021) (cyan). *Right:* azimuthally averaged FWHM of the MeerKAT/Eidos beam normalized by λ/D as a function of frequency (solid cyan line) compared to the ripple model (solid blue line) which is composed of a sinusoidal oscillation on top of a smooth polynomial frequency dependence (solid orange line) - see equation 6. Figures taken from Matshawule et al. (2021).

6.2 Primary beam chromaticity

Measurements of the MeerKAT/Eidos (e.g Asad et al., 2021) beam shows that the FWHM exhibits a low-level frequency-dependent *ripple*. This effect can be seen clearly in the right panel of Figure 4, where the Eidos FWHM is normalized by λ/D . This ripple is caused by the interaction between the primary and secondary reflector of MeerKAT (de Villiers, 2013) and can be important in the foreground cleaning as it will add extra structure to the frequency spectra.

Matshawule et al. (2021) proposed a simple model for the combination of the smooth structure and the ripple shown in the right panel of Figure 4 and valid for the MeerKAT L-band: a high order polynomial with superimposed a sinusoidal oscillation with period T and amplitude A ,

$$\Delta\theta_r = \frac{\lambda}{D} \left(\sum_{d=0}^8 a_d \nu^d + A \sin \left(\frac{2\pi\nu}{T} \right) \right). \quad (6)$$

This type of beam chromaticity complicates component separation appearing as a residual oscillation in the “cleaned” signal. The combined effects of sidelobes and chromaticity has been studied for PCA in Matshawule et al. (2021) in the context of simulations, and the effect can be distinctively seen in the structure of the line-of-sight power spectrum in the left panel of Figure 5. More recently, the effect has also been discussed for other cleaning techniques such as GNILC (see subsection 4.4, De Caro et al. (2025)) and Deep Learning (see subsection 7.3, Ni et al. (2022)).

Emission line stacking analysis of the MeerKLASS L-band deep-field survey confirmed the oscillation structure in observation for the first time (Chen et al., 2025). The chromatic systematics are found to be a convolutional effect on the data, and create correlations in the H_I signal which requires robust covariance estimation to model. Joint inferences on the systematics as well as the H_I signal can help constrain the effect of beam chromaticity.

Note that the ripple in the beam width is expected to be asymmetric between the vertical and the horizontal direction (Asad et al., 2021). The combination of beam asymmetries and sky rotation could result in a superposition of sine waves which will leak the ripple across more scales while possibly reducing its overall amplitude. This can be seen in simulations when a mock HI observation is convolved with Eidos beam maintaining the asymmetries and following a scanning strategy (Matshawule, 2023). The right panel of Figure 5 shows the more complex residual structure found especially in presence of strong point sources.

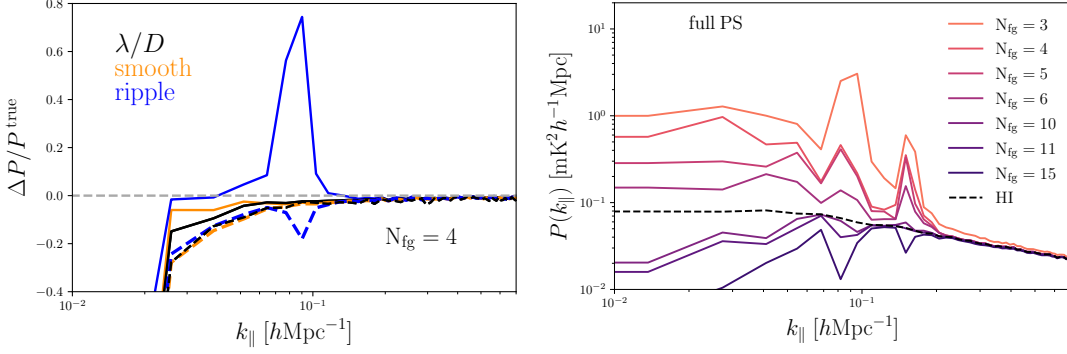


Figure 5: *Left.* Relative difference between the foreground cleaned (with $N_{\text{fg}} = 4$) radial power spectrum and the true input signal for the different FWHM models: $\frac{\lambda}{D}$ in black, the smooth model in orange and the ripple model in blue. Results are shown for the worst and best-case scenario for the point source contamination, considering the full catalogue (solid lines) or only point sources with a flux cut of 100 mJy (dotted lines), respectively. Figure taken from Matshawule et al. (2021). *Right:* Difference, for a mock HI observation with synchrotron and strong point sources, between the cleaned radial power spectrum and the input when also beam asymmetries are considered. Figure taken from Matshawule (2023).

6.3 Joint deconvolution and blind source separation on the sphere (SDecGMCA)

As discussed in the previous section, a frequency-dependent beam can blur the observed sky maps in a way that complicates component separation and applying standard blind component separation methods to these cases can lead to spectral distortions or angular artifacts, ultimately biasing the recovery of the cosmological HI signal. To mitigate these effects, an additional deconvolution step could be incorporated, giving rise to the joint deconvolution and blind source separation problem (DBSS).

In this context, Equation 1 can be extended to account for the convolution with the beam:

$$\mathbf{X}_\nu = (\mathbf{A}_\nu \mathbf{S}) * \mathbf{H}_\nu + \mathbf{N}_\nu, \quad (7)$$

where \mathbf{X}_ν denotes the observed maps at frequency ν , \mathbf{A}_ν is the mixing matrix, \mathbf{S} the matrix of source components, \mathbf{H}_ν the frequency-dependent beam operator, and $*$ denotes convolution.

The presence of a spatially smoothing and frequency-dependent operator—often ill-conditioned or even non-invertible—complicates the separation. The DBSS framework addresses this by jointly estimating both the sources and the mixing matrix while accounting for the beam effects. One such method is SDecGMCA, an extension of the GMCA approach (see subsection 4.3), initially

introduced by Jiang et al. (2017) for sky patches and generalized to full-sky analysis in spherical harmonics by Carloni Gertosio and Bobin (2021)¹.

SDecGMCA reformulates the problem in the spherical harmonic domain, where convolution becomes a simple multiplication:

$$\hat{\mathbf{X}}_v^{\ell,m} = \hat{\mathbf{H}}_v^\ell \mathbf{A}_v \hat{\mathbf{S}}^{\ell,m} + \hat{\mathbf{N}}_v^{\ell,m}, \quad (8)$$

with (ℓ, m) denoting the spherical harmonic coefficients, and where isotropy of the beam implies that $\hat{\mathbf{H}}_v^\ell$ depends only on ℓ .

The algorithm minimizes a cost function that combines data fidelity with a sparsity-promoting prior on the sources:

$$\min_{\mathbf{A}, \hat{\mathbf{S}}} \|\mathbf{\Lambda} \odot (\hat{\mathbf{S}} \mathcal{F}^\dagger \mathbf{\Phi}^\top)\|_1 + \sum_{(\ell,m) \in \mathcal{D}} \|\hat{\mathbf{X}}^{\ell,m} - \text{diag}(\hat{\mathbf{H}}^\ell) \mathbf{A} \hat{\mathbf{S}}^{\ell,m}\|_2^2, \quad (9)$$

where $\mathbf{\Lambda}$ contains thresholding parameters, \odot denotes the element-wise product, $\mathbf{\Phi}$ is a wavelet transform, and \mathcal{F}^\dagger denotes the inverse spherical harmonic transform (see Gkogkou et al. (2026) for more details).

The iterative algorithm alternates between: i) estimating \mathbf{S} (sources) for fixed \mathbf{A} , using a regularized least-squares solution that enforces sparsity and stabilizes the inversion of the beam; ii) updating \mathbf{A} for fixed \mathbf{S} via least-squares; iii) refining the regularization parameters across iterations to ensure convergence and enhanced separation quality.

To evaluate the applicability of SDecGMCA to recovering the H_I cosmological signal, the method was applied to a full-sky simulated data set in Gkogkou et al. (2026). The simulations include cosmological H_I emission alongside dominant astrophysical foregrounds—such as synchrotron radiation, free-free emission, and unresolved point sources. It also incorporates realistic instrumental noise based on MeerKAT-like specifications. In order to assess the robustness of the method in the presence of frequency-dependent instrumental distortions, three beam models of increasing complexity were tested: a Gaussian-degraded beam with constant resolution across frequencies, a Gaussian-evolving beam with a smooth, frequency-dependent width (often assumed in the studies presented in this chapter), and the Gaussian-oscillating beam with the frequency structure studied in Matshawule et al. (2021).

The results (see Figure 6) indicate that SDecGMCA consistently outperforms both standard blind subtraction techniques and parametric fitting methods in the presence of complex, frequency-dependent beam distortions, particularly under the oscillating beam model. Improved recovery is achieved for both the frequency and angular power spectra of the H_I cosmological signal. Furthermore, several reconstruction artifacts previously reported in (Spinelli et al., 2022), including the systematic offset and the spurious peak in the frequency power spectrum, are substantially reduced or entirely suppressed. This result demonstrates the potential of the joint beam deconvolution and cleaning although some improvement are needed to make the algorithm mature for the data. For example, the method has been tested on full-sky simulated observations and a next step would be to incorporate a masking strategy within the framework.

¹<https://github.com/RCarloniGertosio/SDecGMCA>




Figure 6: Angular (left) and frequency (right) power spectrum of the reconstructed H I IM signal using different methods, each represented by a distinct colour. The bottom panels display the ratio of the reconstructed to the true H I power spectrum for each method.

7 Other approaches to component separation

7.1 Foreground removal via the multi-frequency angular power spectrum

The idea is based on the distinctly different spectral behavior of the 21 cm signal as compared to the foregrounds. The signal is expected to be predominantly localized within a typical $\Delta\nu$ range of 0–0.5 MHz, and its amplitude drops substantially and is close to zero at large $\Delta\nu$ (Bharadwaj et al., 2001; Bharadwaj and Sethi, 2001; Bharadwaj and Ali, 2005). On the other hand, the foregrounds generally vary smoothly and remain correlated over a large $\Delta\nu$ range. Following Ghosh et al. (2011a), it is possible to adopt a two-step procedure to remove the foregrounds from the estimated $C_\ell(\Delta\nu)$. First, one can identify a characteristic length scale $[\Delta\nu]$ such that in the range $\Delta\nu > [\Delta\nu]$ the 21 cm signal is predicted to be negligible, and model $C_\ell(\Delta\nu)$ in this range as a combination of foregrounds $[C_\ell(\Delta\nu)]_{\text{FG}}$ and noise

$$C_\ell(\Delta\nu) = [C_\ell(\Delta\nu)]_{\text{FG}} + [\text{Noise}]. \quad (10)$$

The procedure is then to use the $C_\ell(\Delta\nu)$ measured in the range $\Delta\nu > [\Delta\nu]$ to estimate the foreground model $[C_\ell(\Delta\nu)]_{\text{FG}}$.

In the second step, the foreground model can be extrapolated to predict $[C_\ell(\Delta\nu)]_{\text{FG}}$ in the range $\Delta\nu \leq [\Delta\nu]$ and subtract this out from the measured $C_\ell(\Delta\nu)$. Restricting the subsequent analysis to $\Delta\nu \leq [\Delta\nu]$, we expect the residual

$$[C_\ell(\Delta\nu)]_{\text{res}} = C_\ell(\Delta\nu) - [C_\ell(\Delta\nu)]_{\text{FG}} \quad (11)$$

to be a combination of the 21 cm signal and noise

$$[C_\ell(\Delta\nu)]_{\text{res}} = [C_\ell(\Delta\nu)]_{\text{T}} + [\text{Noise}] \quad (12)$$

and we use this to constrain the 21 cm signal $[C_\ell(\Delta\nu)]_{\text{T}}$.

In case of actual data (such as Elahi et al. 2023, 2024), it has been found that the $\Delta\nu$ dependence exhibits a smooth, slowly varying pattern, superimposed on which we have a combination of rapid




Figure 7: The left and the right panels respectively show (blue solid line) the measured $C_\ell(\Delta\nu)$ and $[C_\ell(\Delta\nu)]_T$ for a representative $\ell = 2740$. The black dashed vertical lines show the chosen characteristic length scale $[\Delta\nu]_{0.1}$, where the amplitude of $[C_\ell(\Delta\nu)]_T$ falls to 10% of the maximum correlation, and therefore there is a negligible amount of the signal in the range $\Delta\nu > [\Delta\nu]_{0.1}$. The red solid curve shows the optimally found $[C_\ell(\Delta\nu)]_{FG}$ which minimizes the residual $[C_\ell(\Delta\nu)]_{res}$ in the range $\Delta\nu < [\Delta\nu]_{0.1}$. The dashed curves and the red solid curve in the right panel shows $[C_\ell(\Delta\nu)]_{TP}$ the best-fit polynomials for $[C_\ell(\Delta\nu)]_T$ for the same values of n .

and slow oscillations. The aim of this foreground removal approach is to use the range $\Delta\nu > [\Delta\nu]$ to model the smooth, slowly varying $\Delta\nu$ dependence. It turns out that modeling the oscillatory components either leads to extrapolation biases or a very large signal loss (Elahi et al., 2024). The foreground removal technique is implemented in two different approaches. In the first approach, polynomial fitting (PF), one can model the foregrounds using an even-order polynomial. We use maximum likelihood to estimate the best-fit polynomial coefficients and their error covariance, and use these to obtain the best-fit foreground and the foreground modeling errors in the range $\Delta\nu \leq [\Delta\nu]$. Note that we have restricted the order of the polynomial to model only the smooth, slowly varying component from the measured $C_\ell(\Delta\nu)$. In the second approach, one can use Gaussian Process Regression (GPR; Rasmussen and Williams 2006), which is non-parametric, to model and predict the foreground. This was first implemented in Elahi et al. (2023), where the reader can find more details. In GPR, we model $[C_\ell(\Delta\nu)]_{FG}$ as a Gaussian Process (GP)

$$[C_\ell(\Delta\nu)]_{FG} \sim \mathcal{GP} [0, k_{fg}(\Delta\nu_m, \Delta\nu_n)] , \quad (13)$$

with a zero mean and covariance function (kernel) $k_{fg}(\Delta\nu_m, \Delta\nu_n)$. The idea is to use the $C_\ell(\Delta\nu)$ measured in the range $\Delta\nu > [\Delta\nu]$ to estimate the kernel, which models the smooth, slowly varying foreground component $[C_\ell(\Delta\nu)]_{FG}$. We subsequently use this kernel to predict $[C_\ell(\Delta\nu)]_{FG}$ in the range $\Delta\nu \leq [\Delta\nu]$. We have tried different possible forms for the kernel, and found that the polynomial kernel

$$k_{fg}(\Delta\nu_m, \Delta\nu_n) = c_1 (\Delta\nu_m \cdot \Delta\nu_n + b)^P \quad (14)$$

is well suited for our analysis. Here, the constants c_1 and b are hyper-parameters whose optimal value have been determined by maximizing the log marginal likelihood of the GP, and P , which is not a hyper-parameter, denotes the order of the polynomial kernel. The polynomial kernel is found to model the overall smooth $\Delta\nu$ dependence expected of the foregrounds and causes a minimal signal loss. We have considered different values of P and found that larger values provide a better

fit to the foregrounds in the range $\Delta\nu > [\Delta\nu]$ but at the cost of a larger signal loss. Considering this, we have used $P = 2$ or 3 which provides reasonable choice for HI IM analysis. The GPR predicts the foregrounds $[C_\ell(\Delta\nu)]_{\text{FG}}$ and also the foreground modeling errors in the range $\Delta\nu \leq [\Delta\nu]$. For both PF and GPR, we add the variance of foreground modeling errors to the estimated system noise variance to estimate the total error variance for $[C_\ell(\Delta\nu)]_{\text{res}}$.

7.2 High dimensional Bayesian modeling

Another alternative to blind foreground removal is to attempt to model the foregrounds and instrumental response directly, in order to subtract the non-cosmological components. The high dynamic range between the foregrounds and 21 cm signal makes this challenging, as does the general complexity of instrumental models, and incompleteness of sky models. Any inaccuracy in the sky or instrument model will result in a model error, with even sub-percent errors able to swamp the cosmological signal. Bayesian statistical methods, which not only estimate the best-fit model parameters from the data, but also describe their statistical distribution, can be used to handle the model uncertainty. The uncertainty on the the foreground and systematics can then be propagated into the subsequent analysis steps, mitigating difficulties such as signal loss and other forms of biased results that are present in blind methods. The sky and instrument models need to be flexible enough to incorporate the true unknown behavior of these components. For example: varying source fluxes and power-law spectral indices for the point sources; diffuse foreground emission modeled per-pixel, or in Fourier or spherical harmonic bases. The spectral behavior of the latter can be specified using basis functions such as the principal components found by blind foreground methods, or physics-based approaches such as explicit modeling or moments.

The ever-increasing quantities of data expected for the SKAO necessitates these more sophisticated statistical methods to remain computationally tractable. With Bayesian sampling, models with as few as 100 parameters can begin to suffer from the *curse of dimensionality*: a severe reduction in sampling speed when using more conventional algorithms such as Markov Chain Monte Carlo (MCMC). These performance limitations can be mitigated but not solved using more sophisticated strategies as Hamiltonian Monte Carlo (Duane et al., 1987; Betancourt, 2017, HMC).

One approach which can be resilient to highly-dimensionality is Gibbs sampling, which has been used to great effect in other fields, for example CMB experiments in Eriksen et al. (2008b) which applies the work of Jewell et al. (2004) and Wandelt et al. (2004) to simulated WMAP data, and to observational data in Eriksen et al. (2008a). Further examples include BeyondPlanck Collaboration et al. (2023), which applies Gibbs sampling to observations from the Planck Low Frequency Instrument (LFI); the reprocessing of WMAP data through an extension of the BeyondPlanck framework in Watts et al. (2023); and Watts et al. (2024), which presents the first application of this framework to COBE Diffuse Infrared Background Experiment (DIRBE) data in Cosmoglob DR2. This method allows for the constraint of millions of free model parameters on tractable timescales, assuming some careful modeling choices are made. Rather than sampling from the model's full posterior space, Gibbs sampling instead samples from subsets of this space. By choosing linear models, and ensuring Gaussian priors, these conditional distributions (Lynch, 2007) (or *Gaussian Constrained Realisations*, GCR) can be sampled from iteratively, which is more efficient than sampling from the full posterior. Furthermore, as result of the model's ensured linearity, solutions

can be found through the direct solutions of linear equations, with no need to evaluate likelihood functions.

There are numerous possible applications for Bayesian methods such as this to future SKA interferometric observations for the epoch of reionization (e.g. [Glasscock et al., 2024](#); [Kennedy et al., 2023](#)). An application which specifically addresses single-dish IM can be found in ([Murphy et al., 2026](#)) where the authors achieve a jointly model the 21 cm signal and foregrounds on a timescale of < 1 hour, despite the model containing over 3 million free parameters.

7.3 Deep Learning

In recent years, machine learning techniques have been proposed by many studies as a novel approach for foreground removal. To some degree, subtracting foreground residuals from HI maps is similar to image denoising in the deep learning field. For example, [Makinen et al. \(2021\)](#) applied a deep convolutional neural network (CNN) with a U-Net architecture to simulated single dish data. Compared with PCA, they showed that U-Net can effectively subtract foregrounds from the signals with improved accuracy at intermediate radial scale ($0.025 < k_{\parallel} < 0.075 \text{ h Mpc}^{-1}$). [Ni et al. \(2022\)](#) tested U-Net on simulated MeerKLASS data and found that the U-Net can effectively clean the foreground with both Gaussian and Cosine beam effects. Following [Ni et al. \(2022\)](#), [Gao et al. \(2023\)](#) demonstrated that the U-Net is able to eliminate polarization leakage effects while subtracting 21 cm signal from foregrounds.

While deep learning shows promise for foreground removal in simulations, caution is warranted when applying these methods to real data due to limitations in reliability and generalization. The effectiveness of deep learning models is highly dependent on the training data, which is predominantly based on simulations. However, discrepancies between simulated and real data, stemming from differences in sky model assumptions or instrumental systematics, can compromise the accuracy of foreground subtraction. For example, [Chen et al. \(2024\)](#) evaluated the performance of a U-Net using various sky simulation models for SKA-Mid Band 1. They found that the U-Net performs comparably to PCA only when the training and the test dataset originate consistently from the same sky model. When different models were used for training and testing, the network failed to accurately subtract foregrounds. Moreover, their study demonstrated that the U-Net's ability to handle systematics such as frequency-dependent Gaussian beams and bandpass fluctuation, depends heavily on the inclusion of these systematics in the training phase. [Chen et al. \(2024\)](#) cautions that applying U-Net to real data requires care in two respects: simulated training data may not fully capture the complexity of observations, and the network relies on prior knowledge of systematics being included during training for accurate performance.

Nevertheless, deep learning offers an independent approach to 21 cm foreground removal which, when applied appropriately, can yield complementary results that serve as a valuable cross-check against traditional methods. Currently, deep learning techniques for 21 cm foreground removal often rely on a pre-processing step using traditional methods to reduce the large dynamic range between foregrounds and the signal prior to refined signal reconstruction. Future developments may focus on training the network to directly handle the foreground-to-signal ratio, potentially eliminating the need for pre-processing. Additionally, integrating physical priors, such as the expected power

spectrum of the 21 cm signal or foregrounds, into the network architecture might enhance the model's robustness to variations in real data.

8 Signal loss

Whilst there is a large effort to limit foreground residuals entering \mathbf{N} from Equation 1, causing *additive* contamination, an equally important concern is H_I signal leakage into the removed \mathbf{AS} term, which causes a suppression in the measured H_I power spectrum, known as *signal loss*.

8.1 Transfer function

Constructing a foreground transfer function $\mathcal{T}(\mathbf{k})$ is seen as a promising method for correcting signal loss. This quantifies how the true H_I power spectrum is modulated by the cleaning algorithm. It is defined such that $P_{\text{clean}}(\mathbf{k}) = \mathcal{T}(\mathbf{k}) P_{\text{true}}(\mathbf{k})$, and if an effective means can be devised to provide a transfer function estimate, $\widehat{\mathcal{T}}(\mathbf{k})$, then the signal can be corrected for at the power spectrum level, i.e. this is not a direct reconstruction of the map's attenuation. The unbiased power spectrum used for cosmological inference is thus simply $P_{\text{rec}}(\mathbf{k}) = P_{\text{clean}}(\mathbf{k})/\widehat{\mathcal{T}}(\mathbf{k})$.

In practice, $\widehat{\mathcal{T}}(\mathbf{k})$ is obtained empirically by injecting mock maps of H_I into the real contaminated data, repeating the same blind foreground cleaning, and measuring the fractional suppression of the recovered mock power (Switzer et al., 2015; Cunnington et al., 2023). Using the true data ensures that instrumental effects, noise correlations, and realistic foreground complexity are faithfully represented. The ratio of cleaned to input mock power spectra provides the estimate of the transfer function i.e.

$$\widehat{\mathcal{T}}(k) = \left\langle \frac{\mathcal{P}(\mathbf{M}_{\text{clean}}, \mathbf{M})}{\mathcal{P}(\mathbf{M}, \mathbf{M})} \right\rangle_{N_{\text{mock}}}, \quad (15)$$

where $\mathbf{M}_{\text{clean}}$ is the cleaned mock and real data combination, and \mathbf{M} is the original injected mock, uncleaned. This is averaged over multiple mocks (N_{mock}) to yield a smooth, robust estimate of $\widehat{\mathcal{T}}(\mathbf{k})$ that captures both the scale dependence and anisotropy of the signal loss.

Tests with simulated observations in Cunnington et al. (2023) demonstrated the accuracy of this approach: even when a large number of foreground modes are removed, the reconstructed power spectra remain unbiased at the sub-percent level across the scales well sampled by the survey. These results validate the empirical transfer function correction as a practical and reliable procedure for recovering the true H_I clustering signal from aggressively cleaned data.

Further work in Chen (2025), found that the transfer function correction of signal loss is effectively a normalization of the power spectrum window function in the quadratic estimator formalism (see e.g. Kern and Liu 2021). Through the window function, the effects of foreground cleaning can be understood as a signal attenuation as well as *mode-mixing*. Robust inference using the H_I intensity mapping in the future thus requires accurate modelling of both effects in the power spectrum estimation as well as in the covariance.

8.2 Parametrizing signal loss

As an alternative to correcting for signal loss at the data level through a transfer function, it is possible to incorporate a phenomenological treatment into the model for the comparison with the

measurements and introduce one (or more) nuisance parameters in the fit to constrain the impact of signal loss. Based on simulations, the transfer functions can be well described by a power law behavior: the most straightforward choice is therefore to include a kernel \mathcal{D}_{sl} with such a functional form within the model. In the most general expression, the kernel reads

$$\mathcal{D}_{sl}(k) = \left(\frac{k}{k_{\text{pivot}}} \right)^\beta \quad (16)$$

where k_{pivot} acts as a normalization point and β , that can be left as free parameter, controls the scaling and strength of the effect. More complicated expressions, as broken power laws or kernels that separate the radial and transverse directions, can be applied as well. When including this correction, a generic model of the H_I power spectrum $P_{\text{H}_I}^{\text{th}}(k)$ is modified as $P_{\text{H}_I}^{\text{th}}(k) \rightarrow \mathcal{D}_{sl}(k)P_{\text{H}_I}^{\text{th}}(k)$. This parameterisation is capable of capturing possible contributions from residual foregrounds, although these are expected to be more under control: a positive value of β corresponds to signal loss on large scales, with the severity increasing with β , whereas a negative value indicates an excess of signal. The main limitations of this approach are the possible degeneracies with other parameters of interest (see e.g. [Cunnington et al., 2020](#); [Soares et al., 2021](#)) and the reduced constraining power resulting from the enlargement of the parameter space. On the other hand, it provides a viable alternative to reconstructing the signal itself and it enables a direct estimation of the impact of signal loss from the fits alone, without the need for simulations. This makes it particularly well suited for analyses characterized by a relatively small set of free parameters.

9 Radio Frequency Interference

Radio frequency interference (RFI) is emitted from various terrestrial and orbital sources, such as Global System for Mobile Communication (GSM) antenna tower or satellites and it is one of the biggest issue with an high dynamic range experiment like HiIM. Single-dish observation are particularly sensitive to RFI due to their wide primary beam and sensitivity. Even when observing far way from the main RFI frequency, secondary harmonics may still leak into the data ([Harper and Dickinson, 2018](#)). Moreover, low level broad-band RFI is particularly difficult to remove. As the RFI environment continues to become more challenging, it is crucial in radio observation to find novel algorithms to recognise and possibly mitigate this contamination.

[Engelbrecht et al. \(2025\)](#) simulated RFI from the Radio Navigational Satellite System in MeerKAT L-band in terms of positioning and power seen through the telescope beam, and calibrated the simulation using data. Thanks to the calibration step, they extrapolated the possible contamination out of band, especially for the lower frequencies of interest for cosmology, finding that it seems possible to remove the potential bias with standard component separation methods.

Even when recognised and flagged, RFI pose a severe challenge as it will cause missing data. It is fundamental to test the robustness of the cleaning step and the summary statistic estimation to this problem. [Carucci et al. \(2020\)](#) have found that cleaning methods can deal with missing bands and the power spectrum estimation procedure described in [Cunnington et al. \(2023\)](#) seems also to be robust to missing data. Another way to deal with RFI flags is also to use the multi-frequency angular power spectrum described in [subsection 7.1](#) which easily deals with the missing frequency

channels in the visibility data, which are flagged due to RFI. (Bharadwaj et al., 2019; Pal et al., 2021). One can indeed estimate the MAPS $C_\ell(\Delta\nu)$ first, and then Fourier transform it to estimate the power spectrum $P(k_\perp, k_\parallel)$. In this approach, one can use only the available frequency channels, and it is not essential to compensate for the missing frequency channels.

10 1/f noise and its inclusion in map-making

1/f, pink or flicker noise is a systematic noise bias most notable within single dish data. It is caused by receiver gain fluctuations which, instead of being random are in fact correlated over time and result in a correlated multiplicative effect on the time-ordered-data which can sometimes be seen in the map domain as streaks which follow the path of the scanning strategy. For this reason it is also called stochastic gain or correlated noise. In subsection 10.1, we discuss the modelling of 1/f and its effect of cleaning in the context of single-dish observations and map-making pipelines.

The interplay between smooth and stochastic gain fluctuations, radiometer noise, and sky signal estimation forms, however, a tightly coupled inference problem in autocorrelation measurements. Treating these components in isolation, as discussed in subsection 10.1, can obscure important degeneracies and leads to systematic effects in high dynamic range observations. A recently proposed solution to this is presented in subsection 10.2.

10.1 The 1/f in IM single-dish observations

For measurements taken at a single frequency 1/f noise is parametrized as follows:

$$S(f) = \sigma_n^2 \left(1 + \left(\frac{f_k}{f} \right)^\alpha \right), \quad (17)$$

where σ_n is the Gaussian instrumental noise, f is the Fourier pair of the time elapsed across the time-ordered-data and the knee frequency (f_k) marks the frequency where the 1/f slope, which has a gradient of α , intersects with the thermal noise level (Bigot-Sazy et al., 2015). For a single experiment which measures time-ordered-data across multiple frequency channels, noise can be correlated across both time and frequency and so Harper et al. (2018) considered the following extension to the single channel 1/f noise:

$$\mathbf{S}(f, \tau) = \sigma_n^2 \left(1 + \frac{1}{K\delta\nu} \left(\frac{f_k}{f} \right)^\alpha \left(\frac{\tau_0}{\tau} \right)^{\frac{1-\beta}{\beta}} \right), \text{ where } K = \int d\tau \text{sinc}^2(\pi\delta\nu\tau) \left(\frac{\tau_0}{\tau} \right)^{(1-\beta)/\beta}, \quad (18)$$

with $\tau_0 = (N_\nu\delta\nu)^{-1}$ and N_ν representing the total number of frequency channels. For this 2D parameterization it is the β term which quantifies the frequency correlation; β would be equal to 1 for an experiment where there were no correlations at all between the 1/f noise in each frequency channel, while the closer β tends to zero the stronger the frequency correlations. Mitigating the 1/f noise bias is an important data reduction task for single-dish experiments and has been considered in the context of HI intensity mapping for the SKA-Mid (Chen et al., 2020), FAST (Hu et al., 2021) and BINGO (Wuensche et al., 2022).

Li et al. (2021) placed estimates on the level of 1/f noise within the MeerKLASS survey data through dedicated MeerKAT observations of the South Celestial Pole (SCP). They found that the




Figure 8: Spherically averaged power spectra for simulated $1/f$ noise above the instrumental noise level. Two different scan strategies are considered: HRD and random. Forming sky maps from the time-ordered data was done with (MM) and without (Binned) the inclusion of the noise covariance matrix. Figure adapted from Irfan et al. (2024).

contamination from $1/f$ noise, which on average had knee frequencies of around 0.1 Hz, could be reduced using a 2-mode PCA clean of the time-ordered-data down to the level of $f_k = 0.003$ MHz.

Simulation work in the context of single-dish IM with MeerKAT further investigated the impact of using different observing strategies and map-making algorithms on the reduction of $1/f$ noise. Irfan et al. (2024) measured α , β and f_k parameters from survey data and found the raw $1/f$ noise to be between 10 and 20 times higher than the instrumental noise level and partially correlated over frequency. This partial correlation is what enables successful cleaning (40 per cent reduction in $1/f$ noise power) through the removal of 3 PCA modes. Having experimented with both a random scan strategy and the Horizontal Raster Drift (HRD), as well as map-making assuming white noise properties versus map-making including a noise correlation matrix, they found the HRD strategy optimal for minimizing time-correlated noise and that noise-informed map-making reduced $1/f$ contamination by up to 30 per cent. Figure 8 shows the spherically averaged power spectra of simulated $1/f$ noise above the simulated instrumental noise level. The HRD strategy can be seen to reduce $1/f$ noise compared to the random strategy over all scales, except at around 0.04 Mpc^{-1} . Map-making with a noise covariance matrix (‘MM’) can also be seen to bring down the $1/f$ contamination compared to an assumption an Gaussian noise (‘Binned’).

10.2 Joint inference of $1/f$ noise, gain, and sky signal

Zhang et al. (2026) presents a unified approach that jointly estimates the sky map, the smooth gain components, and the $1/f$ noise parameters, based on a physically motivated data model and performed via high-dimensional Bayesian inference. In this framework, the time-ordered data (TOD) from an autocorrelation detector is modeled as a multiplicative expression of the system temperature, the gain (both smooth and stochastic components), and the radiometer noise. This formulation reflects the radiometer equation in its most general form and naturally encodes the multiplicative coupling between gain, system temperature and noise. By sampling their joint posterior distribution, the method allows information to propagate throughout the inference chain. In particular, the use of time-domain modeling for $1/f$ gain fluctuations avoids the spurious long-range temporal correlations assumed by conventional DFT diagonal $1/f$ models.

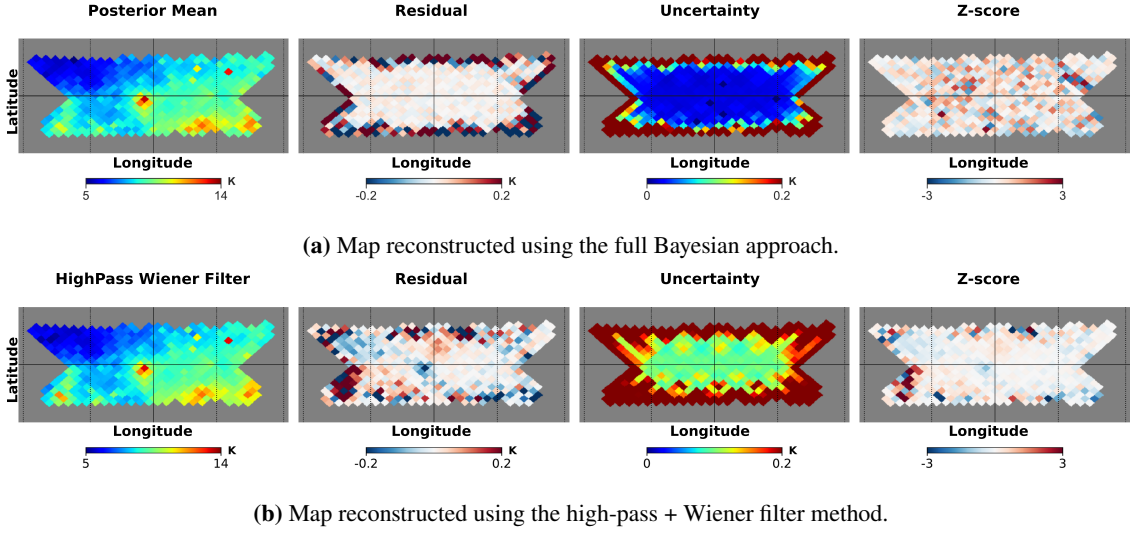


Figure 9: Comparison between the full Bayesian approach and a conventional ‘high-pass filter + Wiener filter’ method. **Residual map:** Map showing the residual difference between the estimated sky signal and the true sky input used in simulations. **Uncertainty map:** Posterior standard deviation of sky parameters marginalized over other parameters, calculated as the sample standard deviation of the draws. **Z-score map:** defined as the residual map divided by the posterior standard deviation. Figure from (Zhang et al., 2026).

To disentangle smooth gain variations from $1/f$ noise, we model the former as a low-order polynomial and the latter as a stationary Gaussian process with a power-law spectrum. The parameter blocks - sky, gain, and noise - are sampled iteratively within a Gibbs framework. Both sky and gain parameters are updated via iterative generalized least squares, which incorporates an effective additive noise estimate. Crucially, the noise amplitude is not fixed or separately parameterized, but intrinsically coupled to the system temperature through gain amplification, ensuring consistent uncertainty propagation. For the efficient sampling of $1/f$ noise parameters, a Levinson-recursive algorithm is employed, which exploits the Toeplitz structure of the time-domain covariance matrix. This reduces computational complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(N^2)$. Using MOMENTEMU (see Zhang 2026) further speeds up the calculation of the noise covariance matrix, enabling good scalability to SKA-scale datasets. The combination of Gibbs sampling techniques and efficient linear solvers for each component ensures computational tractability despite the high dimensionality of the joint inference problem.

To demonstrate the advantages of the full Bayesian approach, Figure 9 compares it with a simple, conventional method that uses a ‘high-pass filter + Wiener filter’ scheme. It can be seen that the traditional method produces larger residuals, particularly within the large-scale structures of the central survey region, and a noisier uncertainty map. Within the survey area, the Bayesian method consistently demonstrates low bias and reduced uncertainty.

11 Discussion and outlook for the SKA-Mid era

In this chapter, we have presented a brief summary of recent advances in component separation and map-making for intensity mapping. The novel methodologies are mostly presented in the context of simulations where it is possible to mimic and isolate specific observational effects. The discussion was necessarily selective, and readers seeking further details are encouraged to consult the references provided.

We have given a large space to blind and semi-blind methods that have been the most successful on data in the last decades (see for example [Cunnington&Wang et al. \(2026\)](#); [Elahi et al. \(2026\)](#)). This trend highlights a difference with the CMB community, where also non-blind and parametric foreground subtraction have always been considered valid and reliable for the extraction of the cosmological information. The reason for such a difference is easily found: the difference in amplitude between the foregrounds and the signal, and the structure of the signal in frequency space, falsifies some of the classical assumptions of component separation in CMB. This trend could change in the next years due to the signal-to-noise improvement expected in the SKAO era.

Whatever the trend may be, the control of subtle instrumental systematics and residual errors from the calibration pipelines will be critical to the success of future Hi IM survey and key to corroborate the solidity of the achievements of IM as a cosmological probe. In conclusion of this chapter, we present a list of the possible challenges left to address.

Ground pickup. Although not an instrumental systematic per se, it does contribute to the temperature budget of single dish experiments (as much as 5K for SKA-Mid Band 1). Its frequency dependence should be smooth; thus, it could be treated as another foreground and can be mitigated keeping the elevation constant. However, its convolution with the beam can create a spectral structure.

Non-linearities. Strong RFI can drive the system non-linear, translating into extra gain fluctuations in time for the autocorrelation signal. If we cannot calibrate fast enough in time, this will generate extra fluctuations in time that will make it harder to remove components such as ground pick up and receiver temperature which do show non-trivial frequency dependences.

Correlated noise. As discussed, $1/f$ can be problematic for single dish experiments. Experience with MeerKAT has shown this to be small for scales of several minutes and also reasonably smooth in frequency, but it will be critical to also check this stability of the SKA-Mid Band 1 receivers.

Frequency bandpass errors. Bandpasses can be well calibrated using good sky models (such as known point sources) but possible changes with time might add frequency fluctuations depending on how often the calibration is repeated. One source of bandpass fluctuations arises from standing waves originating from reflections in the telescope optics. It is common for the optics to change with observing elevation leading to complications for calibration.

The impact of some of these effects, to which single-dish IM could be particularly sensitive, will also depend on the stability and performance of the SKA-Mid. This makes IM a great candidate for the Science Verification. The methodologies developed thanks to instrumental knowledge, or to the discovery of new systematics in real data, allow not only to tackle the immediate challenges that SKAO precursors and pathfinders are facing, but also to pave the way for the SKAO era.

Author List Ordering

The construction of this chapter was led by the first author. All other authors contributed to the writing of the chapter and are listed alphabetically.

Acknowledgements

M. S. is supported by the French government through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence Université Côte d’Azur under reference number ANR- 15-IDEX-01 and by the French Programme National de Cosmologie et Galaxies (PNCG project CIMES). JLB acknowledges funding from the project UC-LIME (PID2022-140670NA-I00), financed by MCIN/AEI/ 10.13039/501100011033/FEDER, UE. IPC is supported by the European Union within the Next Generation EU programme [PNRR-4-2-1.2 project No. SOE_0000136, RadioGaGa]. JF thanks the support of FCT - Fundação para a Ciência e a Tecnologia through national funds by these grants: UID/04434/2025 (DOI 10.54499/UID/04434/2025) and 2023.15069.PEX. JF acknowledges the support from FCT in the form of work through the Scientific Employment Incentive program (reference 2020.02633.CEECIND/CP1631/CT0002) ZZ acknowledges support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 948764) and the RadioForegroundsPlus project HORIZON-CL4-2023-SPACE-01, GA 101135036. JW acknowledges support from the National Natural Science Foundation of China (NSFC, Grant No. 12573112). DC acknowledges support from grant number 192243 of the Swiss National Science Foundation. MGS acknowledges support from the South African Radio Astronomy Observatory and National Research Foundation (Grant No. 84156)

References

- H. Akaike. *IEEE Transactions on Automatic Control*, 19:716–723, Jan. 1974.
- D. Alonso, P. G. Ferreira, and M. G. Santos. *MNRAS*, 444(4):3183–3197, Nov. 2014. doi: 10.1093/mnras/stu1666.
- R. Ansari et al. *AAP*, 540:A129, Apr. 2012. doi: 10.1051/0004-6361/201117837.
- K. M. B. Asad et al. *MNRAS*, 502(2):2970–2983, Apr. 2021. doi: 10.1093/mnras/stab104.
- J. S. Bagla, N. Khandai, and K. K. Datta. *MNRAS*, 407(1):567–580, Sept. 2010. doi: 10.1111/j.1365-2966.2010.16933.x.
- P. Baldi, G. Kerkycharian, D. Marinucci, and D. Picard. *Ann. Statist.*, 37(3):1150–1171, 2009. ISSN 0090-5364,2168-8966. doi: 10.1214/08-AOS601. URL <https://doi.org/10.1214/08-AOS601>.
- R. A. Battye et al. *MNRAS*, 434(2):1239–1256, Sept. 2013. doi: 10.1093/mnras/stt1082.
- C. L. Bennett et al. *APJS*, 148(1):97–117, Sept. 2003. doi: 10.1086/377252.
- M. Betancourt. *arXiv e-prints*, art. arXiv:1701.02434, Jan. 2017. doi: 10.48550/arXiv.1701.02434.

- BeyondPlanck Collaboration et al. *AAP*, 675:A1, July 2023. doi: 10.1051/0004-6361/202244953.
- S. Bharadwaj and S. S. Ali. *MNRAS*, 356:1519–1528, Feb. 2005. doi: 10.1111/j.1365-2966.2004.08604.x.
- S. Bharadwaj and S. K. Sethi. *J. Astrophys. Astron.*, 22:293–307, Dec. 2001. doi: 10.1007/BF02702273.
- S. Bharadwaj, B. B. Nath, and S. K. Sethi. *J. Astrophys. Astron.*, 22:21, Mar. 2001. doi: 10.1007/BF02933588.
- S. Bharadwaj, S. Pal, S. Choudhuri, and P. Dutta. *MNRAS*, 483(4):5694–5700, Mar. 2019. doi: 10.1093/mnras/sty3501.
- M. A. Bigot-Sazy et al. *MNRAS*, 454(3):3240–3253, Dec. 2015. doi: 10.1093/mnras/stv2153.
- J. Bobin, J.-L. Starck, J. Fadili, and Y. Moudden. *IEEE Transactions on Image Processing*, 16(11): 2662–2674, Nov 2007. doi: 10.1109/TIP.2007.906256.
- A. Bonaldi and R. Braun. Square kilometre array science data challenge 1, 2018. URL <https://arxiv.org/abs/1811.10454>.
- A. Bonaldi et al. *Monthly Notices of the Royal Astronomical Society*, 500(3):3821–3837, Oct. 2020. ISSN 1365-2966. doi: 10.1093/mnras/staa3023. URL <http://dx.doi.org/10.1093/mnras/staa3023>.
- J. D. Bowman, M. F. Morales, and J. N. Hewitt. *APJ*, 695(1):183–199, Apr. 2009. doi: 10.1088/0004-637X/695/1/183.
- R. Braun et al. Anticipated performance of the square kilometre array – phase 1 (ska1), 2019. URL <https://arxiv.org/abs/1912.12699>.
- S. Camera et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Camera01.
- R. Carloni Gertosio and J. Bobin. *Digital Signal Processing*, 110:102946, Mar. 2021. doi: 10.1016/j.dsp.2020.102946.
- I. P. Carucci, M. O. Irfan, and J. Bobin. *MNRAS*, Sept. 2020. doi: 10.1093/mnras/staa2854.
- I. P. Carucci et al. *Astron. Astrophys.*, 703:A222, 2025. doi: 10.1051/0004-6361/202453461.
- E. Chapman et al. *MNRAS*, 423:2518–2532, 2012. doi: 10.1111/j.1365-2966.2012.21065.x.
- T. Chen et al. *MNRAS*, 491(3):4254–4266, Jan. 2020. doi: 10.1093/mnras/stz3307.
- T. Chen et al. *Monthly Notices of the Royal Astronomical Society*, 532(2):2615–2634, 07 2024. ISSN 0035-8711. doi: 10.1093/mnras/stae1676. URL <https://doi.org/10.1093/mnras/stae1676>.
- Z. Chen. *MNRAS*, 542(1):L1–L6, Sept. 2025. doi: 10.1093/mnras/slaf058.
- Z. Chen et al. *APJS*, 279(1):19, July 2025. doi: 10.3847/1538-4365/add897.

- S. Choudhuri, S. Bharadwaj, A. Ghosh, and S. S. Ali. *MNRAS*, 445(4):4351–4365, 11 2014. ISSN 0035-8711. doi: 10.1093/mnras/stu2027. URL <https://doi.org/10.1093/mnras/stu2027>.
- S. Choudhuri et al. *MNRAS*, 459(1):151–156, 03 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw607. URL <https://doi.org/10.1093/mnras/stw607>.
- J. J. Condon and S. M. Ransom. *Essential Radio Astronomy*. 2016.
- S. Cunnington, L. Wolz, A. Pourtsidou, and D. Bacon. *MNRAS*, 488(4):5452–5472, 2019. doi: 10.1093/mnras/stz1916.
- S. Cunnington, S. Camera, and A. Pourtsidou. *Mon. Not. Roy. Astron. Soc.*, 499(3):4054–4067, 2020. doi: 10.1093/mnras/staa2986.
- S. Cunnington et al. *MNRAS*, 504(1):208–227, June 2021. doi: 10.1093/mnras/stab856.
- S. Cunnington et al. *Mon. Not. Roy. Astron. Soc.*, 523(2):2453–2477, 2023. doi: 10.1093/mnras/stad1567.
- S. Cunnington&Wang et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Cunnington01.
- B. De Caro et al. *arXiv e-prints*, art. arXiv:2509.02644, Sept. 2025. doi: 10.48550/arXiv.2509.02644.
- D. I. L. de Villiers. *IEEE Transactions on Antennas and Propagation*, 61(5):2457–2465, 2013.
- J. Delabrouille et al. *AAP*, 553:A96, May 2013. doi: 10.1051/0004-6361/201220019. URL <http://adsabs.harvard.edu/abs/2013A%26A...553A..96D>.
- S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. *Physics Letters B*, 195(2):216–222, Sept. 1987. doi: 10.1016/0370-2693(87)91197-X.
- K. M. A. Elahi et al. *MNRAS*, 525(3):3439–3454, Aug. 2023. doi: 10.1093/mnras/stad2495.
- K. M. A. Elahi et al. *MNRAS*, 529(4):3372–3386, Apr. 2024. doi: 10.1093/mnras/stae740.
- K. M. A. Elahi et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Elahi01.
- B. N. Engelbrecht et al. *MNRAS*, 536(1):1035–1055, Jan. 2025. doi: 10.1093/mnras/stae2649.
- H. K. Eriksen et al. *APJL*, 672(2):L87, Jan. 2008a. doi: 10.1086/526545.
- H. K. Eriksen et al. *APJ*, 676(1):10–32, Mar. 2008b. doi: 10.1086/525277.
- J. Fonseca and M. Liguori. *MNRAS*, 504(1):267–279, June 2021. doi: 10.1093/mnras/stab903.
- J. Fonseca et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Fonseca01.
- K. S. F. Fornazier et al. *AAP*, 664:A18, Aug. 2022. doi: 10.1051/0004-6361/202141707.

- L.-Y. Gao, Y. Li, S. Ni, and X. Zhang. *mnras*, 525(4):5278–5290, sep 2023. doi: 10.1093/mnras/stad2646.
- A. Ghosh, S. Bharadwaj, S. S. Ali, and J. N. Chengalur. *MNRAS*, 411(4):2426–2438, 03 2011a. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.17853.x. URL <https://doi.org/10.1111/j.1365-2966.2010.17853.x>.
- A. Ghosh, S. Bharadwaj, S. S. Ali, and J. N. Chengalur. *MNRAS*, 418(4):2584–2589, 12 2011b. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.19649.x. URL <https://doi.org/10.1111/j.1365-2966.2011.19649.x>.
- A. Gkogkou et al. *AAP*, 706:A3, Jan. 2026. doi: 10.1051/0004-6361/202557229.
- K. A. Glasscock et al. *RAS Techniques and Instruments*, 3(1):607–624, Jan. 2024. doi: 10.1093/rasti/rzae041.
- S. E. Harper and C. Dickinson. *MNRAS*, 479(2):2024–2036, Sept. 2018. doi: 10.1093/mnras/sty1495.
- S. E. Harper et al. *MNRAS*, 478(2):2416–2437, Aug. 2018. doi: 10.1093/mnras/sty1238.
- I. Hothi et al. *MNRAS*, 500(2):2264–2277, Jan. 2021. doi: 10.1093/mnras/staa3446.
- W. Hu et al. *MNRAS*, 508(2):2897–2909, Dec. 2021. doi: 10.1093/mnras/stab2728.
- A. Hyvärinen. *IEEE transactions on neural networks*, 10 3:626–34, 1999.
- M. O. Irfan et al. *MNRAS*, 509(4):4923–4939, Feb. 2022. doi: 10.1093/mnras/stab3346.
- M. O. Irfan et al. *MNRAS*, 527(3):4717–4729, Jan. 2024. doi: 10.1093/mnras/stad3457.
- J. Jewell, S. Levin, and C. H. Anderson. *APJ*, 609(1):1–14, July 2004. doi: 10.1086/383515.
- M. Jiang, J. Bobin, and J.-L. Starck. *arXiv e-prints*, art. arXiv:1703.02650, Mar. 2017. doi: 10.48550/arXiv.1703.02650.
- F. Kennedy et al. *APJS*, 266(2):23, June 2023. doi: 10.3847/1538-4365/acc324.
- N. S. Kern and A. Liu. *MNRAS*, 501(1):1463–1480, Feb. 2021. doi: 10.1093/mnras/staa3736.
- Y. Li et al. *MNRAS*, 501(3):4344–4358, Mar. 2021. doi: 10.1093/mnras/staa3856.
- S. M. Lynch. *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*. Springer, New York, 2007.
- S. Majumdar et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Majumdar01.
- T. L. Makinen, L. Lancaster, F. Villaescusa-Navarro, et al. *jcap*, 2021(04):081, apr 2021. doi: 10.1088/1475-7516/2021/04/081.
- D. Marinucci and G. Peccati. *Random fields on the sphere*, volume 389 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2011. ISBN

- 978-0-521-17561-6. doi: 10.1017/CBO9780511751677. URL <https://doi.org/10.1017/CBO9780511751677>. Representation, limit theorems and cosmological applications.
- D. Marinucci et al. *MNRAS*, 383(2):539–545, 2008. doi: 10.1111/j.1365-2966.2007.12550.x.
- S. D. Matshawule. *HI intensity mapping: Impact of primary beam effects*. PhD dissertation, University of the Western Cape, South Africa, 2023. URL <https://uwcscholar.uwc.ac.za/items/55a98e55-9cd4-4309-8110-061158b69edf>.
- S. D. Matshawule, M. Spinelli, M. G. Santos, and S. Ngobese. *MNRAS*, 506(4):5075–5092, Oct. 2021. doi: 10.1093/mnras/stab1688.
- MeerKLASS Collaboration et al. *MNRAS*, 537(4):3632–3661, Mar. 2025. doi: 10.1093/mnras/staf195.
- G. G. Murphy et al. *MNRAS*, 534(3):2653–2673, Nov. 2024. doi: 10.1093/mnras/stae2242.
- G. G. Murphy et al. *arXiv e-prints*, art. arXiv:2604.26890, Apr. 2026. doi: 10.48550/arXiv.2604.26890.
- F. J. Narcowich, P. Petrushev, and J. D. Ward. *SIAM Journal on Mathematical Analysis*, 38(2): 574–594, 2006. doi: 10.1137/040614359. URL <https://doi.org/10.1137/040614359>.
- S. Ni, Y. Li, L.-Y. Gao, and X. Zhang. *apj*, 934(1):83, jul 2022. doi: 10.3847/1538-4357/ac7a34.
- L. C. Olivari, M. Remazeilles, and C. Dickinson. *MNRAS*, 456(3):2749–2765, Mar. 2016. doi: 10.1093/mnras/stv2884.
- S. Pal, S. Bharadwaj, A. Ghosh, and S. Choudhuri. *MNRAS*, 501(3):3378–3391, Mar. 2021. doi: 10.1093/mnras/staa3831.
- D. Pietrobon, A. Balbi, and D. Marinucci. *PRD*, 74(4):043524, Aug. 2006. doi: 10.1103/PhysRevD.74.043524.
- Planck Collaboration et al. *AAP*, 596:A109, Dec. 2016. doi: 10.1051/0004-6361/201629022.
- C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, USA, Jan. 2006.
- M. Remazeilles, J. Delabrouille, and J.-F. Cardoso. *MNRAS*, 418(1):467–476, Nov. 2011. doi: 10.1111/j.1365-2966.2011.19497.x.
- T. Ronconi et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Ronconi01.
- M. Santos et al. In *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, page 19, Apr. 2015.
- M. G. Santos, A. Cooray, and L. Knox. *APJ*, 625(2):575–587, June 2005. doi: 10.1086/429857.
- M. G. Santos et al. *ArXiv e-prints*, Sept. 2017.
- J. R. Shaw et al. *PRD*, 91(8):083514, Apr. 2015. doi: 10.1103/PhysRevD.91.083514.

- SKA Cosmology SWG et al. *PASA*, 37:e007, Mar. 2020. doi: 10.1017/pasa.2019.51.
- P. S. Soares, S. Cunnington, A. Pourtsidou, and C. Blake. *Mon. Not. Roy. Astron. Soc.*, 502(2): 2549–2564, 2021. doi: 10.1093/mnras/stab027.
- P. S. Soares, C. A. Watkinson, S. Cunnington, and A. Pourtsidou. *arXiv e-prints*, art. arXiv:2105.12665, May 2021.
- M. Spinelli, G. Bernardi, and M. G. Santos. *Monthly Notices of the Royal Astronomical Society*, 479(1):275–283, 2018. doi: 10.1093/mnras/sty1457. URL <http://dx.doi.org/10.1093/mnras/sty1457>.
- M. Spinelli et al. *MNRAS*, 509(2):2048–2074, Jan. 2022. doi: 10.1093/mnras/stab3064.
- J.-L. Starck, J. Fadili, and F. Murtagh. *IEEE Transactions on Image Processing*, 16(2):297–309, Jan. 2007. doi: 10.1109/TIP.2006.887733.
- E. R. Switzer et al. *Astrophys. J.*, 815(1):51, 2015. doi: 10.1088/0004-637X/815/1/51.
- B. Thorne, J. Dunkley, D. Alonso, and S. Næss. *Monthly Notices of the Royal Astronomical Society*, 469(3):2821–2833, May 2017. ISSN 1365-2966. doi: 10.1093/mnras/stx949. URL <http://dx.doi.org/10.1093/mnras/stx949>.
- B. D. Wandelt, D. L. Larson, and A. Lakshminarayanan. *PRD*, 70(8):083511, Oct. 2004. doi: 10.1103/PhysRevD.70.083511.
- J. Wang et al. *MNRAS*, 505(3):3698–3721, Aug. 2021. doi: 10.1093/mnras/stab1365.
- D. J. Watts et al. *AAP*, 679:A143, Nov. 2023. doi: 10.1051/0004-6361/202346414.
- D. J. Watts et al. *arXiv e-prints*, art. arXiv:2408.10952, Aug. 2024. doi: 10.48550/arXiv.2408.10952.
- T. L. Wilson and H. S. Rohlfs K. *Tools of Radio Astronomy*. 2013.
- L. Wolz et al. *MNRAS*, 441(4):3271–3283, 2014. doi: 10.1093/mnras/stu792.
- L. Wolz et al. *Monthly Notices of the Royal Astronomical Society*, 510(3):3495–3511, Dec. 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab3621. URL <http://dx.doi.org/10.1093/mnras/stab3621>.
- L. Wolz et al. In *Advancing Astrophysics with the SKA – II (AASKAII)*. 2026. arXiv search: Report number AASKAII/Wolz01.
- L. Wolz et al. *MNRAS*, 464(4):4938–4949, 2017. doi: 10.1093/mnras/stw2556.
- C. A. Wuensche et al. *AAP*, 664:A15, Aug. 2022. doi: 10.1051/0004-6361/202039962.
- J. S. B. Wyithe and A. Loeb. *MNRAS*, 383(2):606–614, Jan. 2008. doi: 10.1111/j.1365-2966.2007.12568.x.
- E. Yohana et al. *MNRAS*, 504(4):5231–5243, July 2021. doi: 10.1093/mnras/stab1197.

Z. Zhang. *MNRAS*, 545(1):staf2039, Jan. 2026. doi: 10.1093/mnras/staf2039.

Z. Zhang, P. Bull, M. G. Santos, and A. Nasirudin. *RAS Techniques and Instruments*, 5:rzag024, Jan. 2026. doi: 10.1093/rasti/rzag024.