

Foreground removal in HI 21 cm intensity mapping under frequency-dependent beam distortions

Athanasia Gkogkou¹, Victor Bonjean¹, Jean-Luc Starck^{2,1}, Marta Spinelli^{3,4}, and Panagiotis Tsakalides^{5,6}

¹ Institutes of Computer Science and Astrophysics, Foundation for Research and Technology Hellas (FORTH), Greece
e-mail: agkogkou@ia.forth.gr

² Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191, Gif-sur-Yvette, France

³ Observatoire de la Côte d’Azur, Laboratoire Lagrange, Bd de l’Observatoire, CS 34229, F-06304 Nice Cedex 4, France

⁴ Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town 7535, South Africa

⁵ Institute of Computer Science, Foundation for Research and Technology Hellas (FORTH), Greece

⁶ Department of Computer Science, University of Crete, Greece

Received 12/09/2025; Accepted 15/11/2025

ABSTRACT

Context. Neutral hydrogen (HI) intensity mapping with single-dish experiments is a powerful approach for probing cosmology in the post-reionization epoch. However, the presence of bright foregrounds over four orders of magnitude stronger than the HI signal makes its extraction highly challenging. While all methods perform well when assuming a Gaussian beam degraded to the worst resolution, most of them degrade significantly in the presence of a more realistic beam model.

Aims. The complexity introduced by frequency-dependent beam effects motivates the need for methods that explicitly account for the instrument’s response. In this work, we investigate the performance of SDecGMCA. This method extends DecGMCA to spherical data, combining sparse component separation with beam deconvolution. Our goal is to evaluate this method in comparison with established foreground removal techniques, assessing its ability to recover the cosmological HI signal from single-dish intensity mapping observations under varying beam conditions.

Methods. We use simulated HI signals and foregrounds informed by existing observational and theoretical models, covering the frequency ranges relevant to MeerKAT and SKA-Mid. The foreground removal techniques tested fall into two main categories: model-fitting methods (polynomial and parametric) and blind source separation methods (PCA, ICA, GMCA, and SDecGMCA). Their effectiveness is evaluated based on the recovery of the HI angular and frequency power spectra under progressively more realistic beam conditions.

Results. While all methods perform adequately under a uniform degraded beam, SDecGMCA remains robust when frequency-dependent beam distortions are introduced. In the oscillating beam case, SDecGMCA suppresses the spurious spectral peak at $k_\nu \sim 0.3$ and achieves $\lesssim 5\%$ accuracy at intermediate angular scales ($10 < \ell < 200$), outperforming other methods. Furthermore, masking bright Galactic regions significantly improves the recovery of the HI signal, particularly for SDecGMCA, which benefits most from excluding contaminated lines of sight. Beam inversion, however, remains intrinsically unstable beyond $\ell \sim 200$, setting a practical limit on the method.

Conclusions. Our findings highlight the limitations of simple fitting and standard blind source separation (BSS) methods once realistic beam effects are considered, and establish SDecGMCA as a particularly promising approach for future single-dish intensity mapping surveys. Its robustness across beam models, combined with the improvements achievable through masking strategies and forthcoming refinements to its thresholding scheme, suggest that SDecGMCA could provide reliable spherical harmonics reconstructions of the HI power spectrum in upcoming experiments.

Key words. large-scale structure – component separation – HI 21 cm intensity mapping

1. Introduction

achieving higher precision in probing the large-scale structure of the Universe with increasing precision requires observations over wider cosmic volumes. One of the most promising probes for this purpose is the 21 cm emission line of neutral hydrogen (HI), which traces the distribution of matter across cosmic time (Furlanetto et al. 2006; Pritchard & Loeb 2012). Unlike optical tracers, the 21 cm line is not obscured by dust and offers a tomographic view of the Universe, as the observed frequency directly maps to redshift.

Due to the faintness of the HI signal, detecting individual galaxies in 21 cm emission beyond the local Universe remains observationally expensive, requiring prohibitively long integration times (e.g., Fernández et al. 2016; Gogate et al. 2020). HI

intensity mapping (IM, Battye et al. 2004; Santos et al. 2005; Chang et al. 2008; Peterson et al. 2009; Kovetz et al. 2017) circumvents this challenge by measuring the collective 21 cm emission from unresolved sources over large sky areas and broad frequency ranges (Battye et al. 2013; Bull et al. 2015). Although this technique sacrifices angular resolution, it enables fast surveys of cosmological volumes with excellent redshift precision.

Several experiments targeting the post-reionization epoch have made significant progress in detecting the HI signal through intensity mapping. The Green Bank Telescope (GBT) has produced cross-correlation measurements with galaxy surveys at low redshift (Masui et al. 2013; Switzer et al. 2013; Wolz et al. 2017; Wolz et al. 2022), while the Parkes radio telescope has reported similar results using data from HIPASS and other galaxy

catalogs (Anderson et al. 2018; Tramonte & Ma 2020; Li et al. 2021). More recently, the MeerKAT radio telescope (Cunnington et al. 2023; Carucci et al. 2024; MeerKLASS Collaboration et al. 2025; Chen et al. 2025) and the Canadian Hydrogen Intensity Mapping Experiment (CHIME) (Amiri et al. 2023, 2024) have added to the growing list of detections, primarily relying on cross-correlations or stacking techniques. Auto-correlation measurements of the HI power spectrum have also emerged, notably at $z \sim 0.32$ and $z \sim 0.44$ (Paul et al. 2023). Complementary work has placed upper limits on the signal (e.g., Mazumder et al. 2025). Achieving high-precision auto-correlation measurements of the cosmological HI power spectrum with well-controlled systematics remains a key goal for the field.

One of the main challenges in HI intensity mapping is contamination from astrophysical foregrounds. Synchrotron emission from the Milky Way, free-free radiation, and extragalactic point sources outshine the HI signal by up to five orders of magnitude, even at high Galactic latitudes (Santos et al. 2005; Jelić et al. 2008). These contaminants exhibit complex spatial and spectral structure, which already complicates separation from the cosmological signal.

In addition, instrumental systematics, in particular the frequency-dependent primary beam, further complicate the problem. The beam response couples to the foregrounds and introduces artificial spectral features that can mimic the 21 cm signal (Matshawule et al. 2021). The situation can be more severe when the beam shows non-smooth (e.g., oscillatory) frequency dependence, as measured for MeerKAT (Asad et al. 2021). Other departures from the idealized chromatic response, such as those caused by instrumental optics or imperfect beam characterization, have also been shown to introduce systematic distortions in both simulations and real data (Shaw et al. 2015; Chen et al. 2025).

A variety of methods have been developed to address this challenge. Among the most widely used are Blind Source Separation (BSS) techniques, which aim to disentangle the HI signal from foreground contamination without relying heavily on prior models. Techniques such as Principal Component Analysis (PCA), Independent Component Analysis (ICA), FastICA, and Generalized Morphological Component Analysis (GMCA) have shown promising results in both theoretical and simulation-based studies (Liu & Tegmark 2011; Ansari et al. 2012; Chapman et al. 2012; Masui et al. 2013; Switzer et al. 2013; Wolz et al. 2014; Bigot-Sazy et al. 2015; Alonso et al. 2015; Shaw et al. 2015; Olivari et al. 2016; Zhang et al. 2016; Wolz et al. 2017; Carucci et al. 2020; Soares et al. 2022).

These BSS techniques, along with several model-fitting approaches, were systematically compared in Spinelli et al. (2021), where their performance was evaluated under various beam configurations. That study found that no method performs optimally across all regimes. Each technique exhibited specific limitations, including systematic offsets and artifacts arising from beam effects or intrinsic properties of the foreground and HI signals.

With the rise of machine learning, there has been increasing interest in applying deep learning techniques to foreground removal and HI signal reconstruction (Makinen et al. 2021; Ni et al. 2022; Acharya et al. 2023; Shi et al. 2024; Ni et al. 2024; Chen et al. 2024). These approaches have demonstrated notable success in simulation-based studies, often leveraging large datasets such as CAMELS (Villaescusa-Navarro et al. 2021). However, their reliance on high-fidelity training simulations makes them sensitive to mismatches between training and real-world data. Incomplete knowledge of astrophysical foregrounds or instrumental systematics can bias these methods and poten-

tially compromise their robustness. To avoid such dependencies, we opt for a fully blind, model-independent approach in this work.

A promising solution is to jointly address beam deconvolution and source separation, since the standard linear mixture model ($\mathbf{Y} = \mathbf{A}\mathbf{X}$) becomes invalid when the beam differs across frequency channels (Jiang et al. 2017). In such cases, deconvolution and component separation must be performed simultaneously rather than sequentially. The DecGMCA method introduced in that work achieves this by incorporating the concept of sparsity, originally developed in GMCA (Bobin et al. 2007; Starck et al. 2016), while also addressing the deconvolution problem. Its extension to full-sky data, SDecGMCA (Carlioni Gertosio & Bobin 2021), enables joint source separation and beam correction on the sphere. In this work, we apply SDecGMCA for the first time to simulated HI single-dish intensity mapping data at post-reionization frequencies. By revisiting the comparison of widely used foreground removal methods, we evaluate whether SDecGMCA can resolve key limitations reported in earlier analyses.

The paper is structured as follows. In Sect. 2, we describe the simulations used to evaluate the performance of different foreground removal techniques. Sect. 3 describes the new method we propose along with the other component separation algorithms used for comparison. In Sect. 4, we present and compare the performance of these methods across different instrumental configurations. In Sect. 5, we examine how a masking strategy can enhance the performance of the foreground removal methods. Sect. 6 places our findings in the context of previous work and provides a final summary.

2. Simulations

In this section, we describe the simulated data used to test and validate the various foreground removal methods. To create a realistic representation of the sky in the frequency range 900–1400 MHz, the simulations include multiple components: the 21 cm cosmological signal, galactic foregrounds such as synchrotron and free-free diffuse emission, and extragalactic foregrounds.

For each frequency and component, we generate HEALPix maps (Górski et al. 2005) with $N_{\text{side}} = 256$, corresponding to $N_{\text{pix}} = 12 \times N_{\text{side}}^2$ pixels per map. These components are then combined into full-sky maps, convolved with various telescope beam models, and finally contaminated with white noise per frequency channel, calculated using standard thermal noise estimates.

2.1. Astrophysical components

2.1.1. HI cosmological signal

At low redshifts, neutral hydrogen is expected to reside only in high-density regions where it is self-shielded from the ionizing effects of the cosmic ultraviolet background. As a result, HI traces the densest regions of the underlying dark matter field, making it a linearly biased tracer of dark matter. To simulate the HI distribution, we use the CRIME simulation¹ (Alonso et al. 2014). This software approximates the dark matter field using a lognormal realization. In this work the adopted cosmological parameters are $\{\Omega_m, \Omega_\Lambda, \Omega_b, h\} = \{0.3, 0.7, 0.049, 0.67\}$. The initial simulation cube has a side length of $3 \text{ Gpc } h^{-1}$ and

¹ <http://intensitymapping.physics.ox.ac.uk/CRIME.html>

is divided into 2048^3 cells. Light-cone effects and redshift-space distortions are naturally incorporated into the simulation. The HI density field is then obtained by applying an assumed linear HI bias. Following the prescription in [Carucci et al. \(2020\)](#), we adopt a redshift-dependent linear HI bias of the form $b_{\text{HI}}(z) = 0.3(1+z) + 0.6$ ([Martin et al. 2012](#)), which is consistent with observational constraints at redshifts $z \lesssim 0.8$ ([Switzer et al. 2013](#)). For the HI cosmic abundance, we assume a redshift evolution given by $\Omega_{\text{HI}}(z) = 4 \times 10^{-4}(1+z)^{0.6}$, in agreement with measurements from [Crichton et al. \(2015\)](#).

2.1.2. Synchrotron emission

We adopt the synchrotron emission model implemented in the Python Sky Model 3 (PySM3; [Zonca et al. 2021](#)), which describes the emission as a power law in frequency with a spatially varying spectral index. The synchrotron brightness temperature at frequency ν and sky position p is given by:

$$T_{\text{sy}}(\nu, p) = T_s(p) \left(\frac{\nu}{\nu_0} \right)^{\beta_{\text{sy}}(p)}, \quad (1)$$

where $T_s(p)$ is the amplitude map at the reference frequency ν_0 , and $\beta_{\text{sy}}(p)$ is the spectral index map.

The amplitude template is based on the 408 MHz all-sky map reprocessed by [Remazeilles et al. \(2015\)](#), while the spectral index map is derived from a combination of the Haslam 408 MHz data and WMAP 23 GHz observations ([Miville-Deschênes et al. 2008](#)). Small-scale fluctuations have been added directly to the intensity template to better emulate observed structure. This model corresponds to the “power-law” option of the Planck Sky Model v1.7.8 ([Delabrouille et al. 2013](#)), updated with the Remazeilles version of the Haslam map.

2.1.3. Free-free emission

We also use the PySM3 prescription for the free-free emission. The spatial template is based on the analytic model adopted in the Commander fit to the Planck 2015 data, following the formalism described by [Draine \(2011\)](#). This template represents the free-free brightness temperature at 30 GHz on degree angular scales. To represent small-scale fluctuations, additional structure is added following the procedure outlined in [Thorne et al. \(2017\)](#). The native resolution of the map is inherited from PySM2.

The emission is scaled in frequency using a spatially uniform power-law spectral index of $\beta_{\text{ff}} = -2.14$, yielding the following model for the free-free brightness temperature:

$$T_{\text{ff}}(\nu, p) = T_{\text{ff}}(p) \left(\frac{\nu}{\nu_0} \right)^{\beta_{\text{ff}}}, \quad (2)$$

where $T_{\text{ff}}(p)$ is the free-free amplitude map at the reference frequency $\nu_0 = 30$, GHz.

2.1.4. Extragalactic point sources

For extragalactic radio sources, an empirical model is adopted [Battye et al. \(2013\)](#), which fits a fifth-order polynomial to observational source counts at 1.4 GHz from various surveys. By integrating these source counts, we derive the mean temperature representing unresolved sources at 1.4 GHz. The spatial distribution of these sources is characterized by the angular power spectrum, which consists of both clustering and Poisson components.

The clustering term is derived from the observed angular correlation function of radio sources, while the Poisson component is obtained by integrating the differential source counts weighted by their flux squared. The total angular power spectrum is given by:

$$C_\ell = C_\ell^{\text{clust}} + C_\ell^{\text{Poisson}}. \quad (3)$$

It is then transformed into pixel space, mapping the sources by using the HEALPix synfast routine. Point sources exceeding 0.01 Jy are randomly injected as fully resolved sources, following [Olivari et al. \(2016\)](#):

$$T_{\text{ps}}(1.4 \text{ GHz}, p) = \left(\frac{\lambda^2}{2k_B} \right) \Omega_{\text{pixel}}^{-1} \sum_{i=1}^N S_i, \quad (4)$$

where $\lambda^2/2k_B$ is the conversion factor between flux and brightness temperature, Ω_{pixel} is the map pixel area, and N represents the number of sources per steradian with flux S , as determined by the empirical fit of the differential counts. To scale this 1.4 GHz estimate across the requested frequency range, we apply a power law with a spectral index $\beta_{\text{ps}}(p)$ that varies across the sky. This index follows a Gaussian distribution centered at -2.7 with a standard deviation of 0.2 ([Bigot-Sazy et al. 2015](#)). Consequently, the temperature of extragalactic point sources at each pixel and frequency is given by:

$$T_{\text{ps}}(\nu, p) = T_{\text{ps}}(1.4 \text{ GHz}, p) \left(\frac{\nu}{\nu_0} \right)^{\beta_{\text{ps}}(p)}, \quad (5)$$

where $\beta_{\text{ps}}(p) \sim \mathcal{N}(-2.7, 0.2)$. Sources brighter than 1 Jy are assumed to have been identified and removed from the data.

2.2. Instrumental effects

After generating and merging all components, we account for two key instrumental effects in the maps. First, the application of the telescope beam, which smooths the signal. Second, the addition of uncorrelated thermal noise.

2.2.1. Telescope beam models

Our goal is to evaluate the performance of foreground removal methods in the presence of a beam model, progressing from simple to more complex approaches. In this study, we focus on a single-dish experiment with specifications similar to the MeerKAT radio telescope, featuring a dish diameter of 13.5 m. However, our analysis remains broadly applicable to other instrument configurations.

The most straightforward beam model assumes a spherically symmetric Gaussian profile with no side lobes. The beam response is given by:

$$B(\theta) = e^{-\frac{\theta^2}{2\sigma^2}}, \quad (6)$$

where θ represents the angular distance from the beam center, and the standard deviation σ is related to the full width at half maximum (FWHM) by:

$$\sigma = \frac{\text{FWHM}}{\sqrt{8 \ln 2}}. \quad (7)$$

The FWHM of the beam is frequency-dependent and is given by:

$$\text{FWHM} = \frac{\lambda}{D} = \frac{c}{\nu D}, \quad (8)$$

where c is the speed of light, λ the observing wavelength, ν the observing frequency, and D the telescope dish diameter. To achieve a uniform angular resolution across all frequency channels, we degrade the maps to match the resolution of the lowest-frequency channel, corresponding to the largest FWHM. This requires an additional smoothing step, where the standard deviation in Eq. 6 is adjusted using:

$$\sigma_{\text{new}} = \sqrt{\sigma_{\text{worst res}}^2 - \sigma^2}, \quad (9)$$

where $\sigma_{\text{worst res}} = \text{FWHM}_{\text{max}} / \sqrt{8 \ln 2}$ corresponds to the lowest-resolution channel, and σ is the original beam standard deviation for a given frequency.

When this additional smoothing step is applied, we refer to the resulting beam model as the "Gaussian degraded" model. Conversely, if no re-smoothing is performed and the beam varies naturally with frequency, we denote it as the "Gaussian evolving" beam model. These two models will be referenced throughout our analysis.

To achieve a more realistic approximation of the MeerKAT primary beam, we introduce an oscillatory term in Eq. 8 to account for additional effects arising from the interaction between the primary and secondary reflectors, as observed and modeled by Matshawule et al. (2021). Their study compared the simple Gaussian beam assumption to a more accurate representation of the MeerKAT beam and characterized the frequency-dependent evolution of the FWHM using the following expression:

$$\text{FWHM} = \frac{\lambda}{D} \left[\sum_{d=0}^8 a_d \nu^d + A \sin\left(\frac{2\pi\nu}{T}\right) \right], \quad (10)$$

where the first term is an 8th-degree polynomial (with coefficients provided in Table 1 of Matshawule et al. 2021), and the second term introduces a sinusoidal oscillation with a period of $T = 20$ MHz and an amplitude of $A = 0.1$ arcmin. Throughout this work, we refer to this refined beam model as the "Gaussian oscillating" beam model.

2.2.2. Instrumental noise

For this study, we adopt instrument and survey parameters representative of an SKA1-MID-like single-dish survey (Santos et al. 2015). Using Eq. 11, we generate full-sky noise maps, treating the computed values as the per-pixel noise variance at each frequency channel.

The instrumental noise is modeled as Gaussian and uncorrelated, with uniform variance across the sky. Its standard deviation varies with frequency according to

$$\sigma_N(\nu) = T_{\text{sys}}(\nu) \sqrt{\frac{4\pi f_{\text{sky}}}{\Delta\nu t_{\text{obs}} N_{\text{dishes}} \Omega_{\text{beam}}(\nu)}}, \quad (11)$$

where T_{sys} is the system temperature, f_{sky} the observed sky fraction, $\Delta\nu$ the channel width, t_{obs} the total observing time, N_{dishes} the number of telescope dishes, and Ω_{beam} the beam solid angle. The latter is related to the full width at half maximum (FWHM) of the beam through $\Omega_{\text{beam}}(\nu) = 1.133 \theta_{\text{FWHM}}^2(\nu)$.

The system temperature, which includes contributions from both receiver noise and sky emission, is modeled as

$$T_{\text{sys}}(\nu) = T_{\text{instr}} + 66 \left(\frac{300 \text{ MHz}}{\nu} \right)^{2.55} \text{ K}, \quad (12)$$

where T_{instr} is the instrumental contribution and ν is the observing frequency in MHz.

The adopted survey parameters are: observed sky fraction $f_{\text{sky}} = 1$ (full-sky coverage), channel width $\Delta\nu = 1$ MHz, total observing time $t_{\text{obs}} = 20,000$ h, instrumental temperature $T_{\text{instr}} = 25$ K, and number of dishes $N_{\text{dishes}} = 197$. The relatively long observing time was chosen to minimize noise impact and to focus the comparison on method performance.

3. Methods

An intensity mapping survey produces a map of the total brightness temperature T for each frequency channel ν . At any given sky pixel p , the observed temperature is a combination of three components: the cosmological HI 21 cm signal ($T_{\nu,p}^{\text{HI}}$), astrophysical foregrounds ($T_{\nu,p}^{\text{cont}}$), and instrumental noise ($T_{\nu,p}^{\text{noise}}$). This can be written as:

$$T_{\nu,p}^{\text{obs}} = T_{\nu,p}^{\text{cont}} + T_{\nu,p}^{\text{HI}} + T_{\nu,p}^{\text{noise}}. \quad (13)$$

The objective is to extract the cosmological component $T_{\nu,p}^{\text{HI}}$ from the total observed signal $T_{\nu,p}^{\text{obs}}$. Since the HI signal is mixed with instrumental noise, the separation focuses on removing the dominant foreground component. This can be achieved through two main approaches: (i) fitting and subtracting a smooth continuum model, ideally isolating the residual HI+noise component or (ii) applying blind source separation (BSS) algorithms under the assumption that the continuum emission is sufficiently smooth to be represented by a limited number of components. In this section, we provide a detailed overview of both approaches: model-fitting methods and BSS techniques.

3.1. Blind source separation (BSS)

In many scientific problems, especially in astrophysics, the primary objective is to extract meaningful information from multi-channel observations. BSS, an unsupervised matrix factorization technique, is widely employed for this purpose. Specifically, we consider data consisting of N_c frequency channels, where each channel provides an observation (X_ν). This observation is a linear combination of N_s underlying sources (S), with the addition of some level of noise (N_ν). Therefore, for channel ν , a single observation, which represents an image, can be described by the following linear mixture model:

$$X_\nu = \sum_{n=1}^{N_s} A_{\nu,n} S_n + N_\nu, \quad (14)$$

where A is the mixing matrix, which determines the contribution of each source S_n to the observation at frequency ν . The goal of any BSS method is to estimate both A and S in an unsupervised way, relying solely on the observed data X .

However, this is an ill-posed inverse problem, as multiple combinations of A and S can yield the same observations X . To overcome this ambiguity, additional prior information about the sources or the mixing matrix is required. Commonly used constraints include assuming statistical independence of the sources

(ICA), enforcing non-negativity of both sources and mixing coefficients (non-negative matrix factorization), or leveraging source morphological diversity and sparsity in a specific representation domain (GMCA). The problem becomes even more complex when the data are blurred due to the instrument's point spread function (PSF) or beam. In such cases, beyond source separation, an additional deconvolution step is necessary to recover the true signal (SDecGMCA). In the following, we describe these different approaches in more detail.

3.1.1. Principal component analysis (PCA)

Principal component analysis (PCA, Maćkiewicz & Ratajczak 1993) is a statistical method assumes that the sources are uncorrelated and seeks to diagonalize the covariance matrix of the observations, given by $\mathbf{C} = \mathbf{X}\mathbf{X}^\top$. Achieving this means expressing the covariance matrix in the form:

$$\mathbf{C} = \mathbf{U}\mathbf{D}\mathbf{U}^\top, \quad (15)$$

where $\mathbf{D}$ is a diagonal matrix containing the eigenvalues. Since the observed data $\mathbf{X}$ is a linear combination of the source signals $\mathbf{S}$, weighted by the mixing matrix $\mathbf{A}$ ($\mathbf{X} = \mathbf{A}\mathbf{S}$), we can rewrite the covariance matrix as:

$$\mathbf{C} = (\mathbf{A}\mathbf{S})(\mathbf{A}\mathbf{S})^\top = \mathbf{A}\mathbf{S}\mathbf{S}^\top\mathbf{A}^\top = \mathbf{U}\mathbf{D}\mathbf{U}^\top, \quad (16)$$

bringing it to the form of Eq. 15, where $\mathbf{D} = \mathbf{S}\mathbf{S}^\top$ and $\mathbf{U} = \mathbf{A}$. If PCA successfully diagonalizes $\mathbf{C}$, it effectively determines the mixing matrix $\mathbf{A}$, whose columns correspond to the principal components of the data. These principal components are the eigenvectors of $\mathbf{C}$, computed after centering $\mathbf{X}$ to have zero mean.

Because bright foregrounds dominate both the amplitude and variance of the observed signal, they are expected to be captured by the first few principal components, those corresponding to the largest eigenvalues. This property allows PCA to efficiently isolate and remove foreground contamination, revealing the underlying cosmological 21 cm signal. As a result, PCA has become a widely adopted method for foreground cleaning in the HI intensity mapping community (e.g., Alonso et al. 2015; Cunnington et al. 2021).

3.1.2. Fast independent component analysis (FastICA)

This statistical method, which was developed in Hyvarinen (1999) estimates the mixing matrix $\mathbf{A}$ by leveraging the assumption that sources are statistically independent. FastICA maximizes statistical independence using non-Gaussianity as a proxy. According to the central limit theorem, the probability density function of a mixture of independent variables is more Gaussian than that of any individual variable. Therefore, by maximizing a statistical measure of non-Gaussianity, FastICA effectively isolates the statistically independent sources.

Before applying FastICA, the data must first be centered to zero mean. The method then achieves non-Gaussianity maximization by maximizing the negentropy of the sources, a measure that quantifies the distance of a distribution from Gaussianity, highlighting structured patterns within the sources. FastICA has been successfully applied to both simulated and observational HI intensity mapping data for foreground removal and cosmological signal recovery (e.g., Wolz et al. 2014; Cunnington et al. 2019; Hothi et al. 2021).

3.1.3. Generalized morphological component analysis (GMCA)

This method relies on two key assumptions. First, the sources are sparse, meaning they have mostly zero values in a suitable transformed domain, such as the wavelet domain. Second, the sources show morphological diversity. This means their non-zero values appear in distinct, non-overlapping regions. These assumptions are essential as they allow to dramatically improve the contrast between distinct components and eventually ease the separation process.

This method solves Eq. 14 iteratively after first wavelet transforming the data from $\mathbf{X}$ to $\mathbf{X}^{\text{wt}}$, using the starlet wavelet dictionary (Starck et al. 2007). GMCA aims to minimize the following cost function:

$$\min_{\mathbf{A}, \mathbf{S}^{\text{wt}}} \sum_{i=1}^{N_s} \lambda_i \|\mathbf{S}^{\text{wt}} \mathbf{i}\|_1 + \|\mathbf{X}^{\text{wt}} - \mathbf{A}\mathbf{S}^{\text{wt}}\|_F^2, \quad (17)$$

where the first term represents the sparsity constraint, and the second term is the data fidelity term. Here, $\|\cdot\|_1$ denotes the ℓ_1 norm, defined as $\|\mathbf{M}\|_1 = \sum_{i,j} |m_{i,j}|$, and $\|\cdot\|_F$ is the Frobenius norm, which is defined as $\|\mathbf{M}\|_F^2 = \sum_i \sum_j |m_{i,j}|^2 = \text{Trace}(\mathbf{M}\mathbf{M}^\top)$. λ_i are regularization coefficients, or sparsity thresholds, ensuring that noise is robustly removed. These coefficients are gradually decreased to a final level determined by the noise characteristics (Bobin et al. 2007). GMCA² is a blind method, meaning that no specific model is used for either $\mathbf{A}$ or $\mathbf{S}$. The only required inputs are the number of components, N_s , and the assumption of sparsity.

This method has been previously applied to cosmic microwave background (CMB) map reconstruction using WMAP and Planck data (Bobin et al. 2014, 2016), to 21 cm signal recovery (Chapman et al. 2013; Carucci et al. 2020; Marins et al. 2022), and more recently to the study of supernova remnants in X-ray data (Godinaud et al. 2023).

3.2. Fitting methods

Alternative approaches involve fitting methods, which aim to smooth the signal, effectively filtering out noise and isolating the foregrounds that vary smoothly with frequency. Such methods have been widely used in the literature for foreground cleaning (e.g., Spinelli et al. 2021; Fonseca & Liguori 2021).

3.2.1. Polynomial fitting

The simplest approach to foreground subtraction involves modeling the foregrounds with a smooth frequency-dependent function. We represent this smooth baseline with a 6th-order polynomial for each line of sight (p) as a function of frequency ν :

$$T_{\text{fg}}(\nu, p) = \sum_{n=0}^6 a_n(p) \nu^n. \quad (18)$$

Here, $a_n(p)$ are the polynomial coefficients for each line of sight, capturing the smoothly varying foreground component across frequencies. The polynomial fit is then subtracted from the total temperature maps, effectively flattening the signal and removing the smooth foreground baseline.

However, a second step is required to prevent overfitting bright HI sources, whose distinct profiles may be misinterpreted

² <https://www.cosmostat.org/statistical-methods/gmca>

by the polynomial fit and subsequently removed. To address this, we perform sigma clipping after the initial polynomial subtraction setting the threshold to 5σ . This step identifies and masks outliers, effectively marking the locations of extremely bright sources. With these outliers excluded, we reapply the polynomial fit to capture only the smooth foreground baseline without being influenced by the bright sources.

3.2.2. Parametric fitting

Another approach is to leverage prior empirical knowledge of the foregrounds about their spatial distribution and frequency evolution. Synchrotron, free-free emission, and extragalactic point sources exhibit strong degeneracy due to their similar spectral characteristics (Planck Collaboration et al. 2016). However, we assume that at the frequencies of interest (MHz), the diffuse synchrotron and free-free emissions are the major contributors. Therefore, in this analysis, we only model these two components. We assume that the extragalactic contribution will be absorbed into one or both of these components. Consequently, for our parametric fit, we adopt the model:

$$T(\nu) = \alpha_1(p) \left(\frac{\nu}{\nu_0} \right)^{-2.13} + \alpha_2(p) \left(\frac{\nu}{\nu_0} \right)^{\beta(p)}, \quad (19)$$

where $\alpha_1(p)$, $\alpha_2(p)$, and $\beta(p)$ are free parameters determined for each line of sight p . The first term corresponds to the free-free emission, with the power law fixed at -2.13 (Bennett et al. 1992). The second term accounts for the remaining diffuse emission, i.e., synchrotron, where the power law is treated as a free parameter.

3.3. A new method for beam-aware foreground separation: SDecGMCA

When observations are affected by frequency-dependent distortions, such as a varying telescope beam, standard BSS and model-fitting methods struggle to disentangle beam-induced effects from the intrinsic spectral and spatial properties of the sky components. To address this challenge, we apply SDecGMCA (Carloni Gertosio & Bobin 2021, hereafter C21) for the first time in the context of HI intensity mapping. SDecGMCA extends DecGMCA by performing joint blind source separation and beam deconvolution on full-sky spherical data, incorporating the frequency-dependent beam response directly into the separation model.

In this framework, the standard BSS model is generalized as:

$$\mathbf{X}_\nu = (\mathbf{A}_\nu \mathbf{S}) * \mathbf{H}_\nu + \mathbf{N}_\nu, \quad (20)$$

where $\mathbf{X}_\nu$ is the observed data at frequency ν , $\mathbf{A}_\nu$ is the mixing matrix, $\mathbf{S}$ is the source matrix, $\mathbf{H}_\nu$ is the beam operator (assumed isotropic and thus diagonal in harmonic space), and $*$ denotes spherical convolution. $\mathbf{N}_\nu$ represents additive noise. The model is reformulated in the spherical harmonic domain as:

$$\hat{\mathbf{X}}_\nu^{\ell m} = \hat{\mathbf{H}}_\nu^\ell \mathbf{A}_\nu \hat{\mathbf{S}}^{\ell, m} + \hat{\mathbf{N}}_\nu^{\ell, m}, \quad (21)$$

where ℓ and m are the spherical harmonic indices, and $\hat{\mathbf{H}}_\nu^\ell$ is the beam transfer function at frequency ν . The beam affects only ℓ due to isotropy, simplifying the deconvolution. Throughout this work, boldface symbols (e.g., $\mathbf{X}$, $\mathbf{A}$, $\mathbf{S}$) denote matrices or linear operators, while non-bold symbols (e.g., $\hat{\mathbf{H}}_\nu^\ell$) represent scalar

quantities. A hat (e.g., $\hat{\mathbf{S}}$) indicates that the quantity is expressed in the spherical harmonic domain.

The inverse problem of recovering $\mathbf{A}$ and $\hat{\mathbf{S}}$ is addressed by minimizing the following cost function:

$$\min_{\mathbf{A}, \hat{\mathbf{S}}} \left\| \mathbf{A} \odot (\hat{\mathbf{S}} \mathcal{F}^\dagger \Phi^\top) \right\|_1 + \sum_{(\ell, m) \in \mathcal{D}} \left\| \hat{\mathbf{X}}^{\ell, m} - \text{diag}(\hat{\mathbf{H}}^\ell) \mathbf{A} \hat{\mathbf{S}}^{\ell, m} \right\|_2^2, \quad (22)$$

where Φ is the starlet wavelet transform and $\mathcal{F}^\dagger$ is the inverse spherical harmonic transform. The first term promotes sparsity in the wavelet domain via soft-thresholding, with thresholds controlled by $\mathbf{A}$, while the second term enforces data fidelity in harmonic space.

Iterative algorithm and regularization strategy

SDecGMCA proceeds in two main stages: a warm-up phase and a refinement phase. Each iteration alternates between estimating the source maps $\hat{\mathbf{S}}$ and the mixing matrix $\mathbf{A}$. The main steps and their corresponding hyperparameters are summarized below.

- (i) Initialization of $\mathbf{A}$: The mixing matrix is initialized via a PCA on data re-convolved to the lowest-resolution channel.
- (ii) Update of $\hat{\mathbf{S}}$: The sources are estimated using a Tikhonov-regularized least-squares solution for each (ℓ, m) (Eq. 21 in C21), where the regularization coefficient $\epsilon_{n, \ell}$ is adaptively set according to the selected strategy (Eqs. 24–28 in C21). The overall strength of the regularization is controlled by the hyperparameter c .
- (iii) Enforcing sparsity: After the inverse harmonic transform, $\hat{\mathbf{S}}$ is projected onto an isotropic undecimated wavelet (starlet) basis with five scales. Sparsity is promoted via a soft-thresholding scheme, where the number of retained coefficients (active support) increases linearly from 0 to the hyperparameter $K_{\max}$ during the warm-up phase. During the refinement stage, an ℓ_1 reweighting step is applied to mitigate thresholding bias.
- (iv) Update of $\mathbf{A}$: The mixing matrix is estimated via least squares over all spherical harmonics (Eq. 14 in C21). After each update, the columns of $\mathbf{A}$ are normalized to prevent degeneracies, and negative entries are set to zero to enforce physical non-negativity (Eqs. 15–16 in C21).
- (v) Refinement phase: During refinement, the thresholding levels and regularization parameters are adapted more precisely (see Table 1 in C21) to improve convergence and mitigate overfitting.

Iterations continue until the relative change in the source estimate satisfies $\|\hat{\mathbf{S}}^{(\text{it})} - \hat{\mathbf{S}}^{(\text{it}-1)}\|_2 / \|\hat{\mathbf{S}}^{(\text{it})}\|_2 < \epsilon$, with $\epsilon = 10^{-2}$ during the warm-up phase and $\epsilon = 10^{-4}$ during refinement.

Among the regularization strategies proposed in C21, we adopt Strategy #3 for the warm-up phase. This approach sets the Tikhonov regularization coefficient $\epsilon_{n, \ell}$ based on the inverse of the smallest eigenvalue of the matrix $\mathbf{M} = \mathbf{A}^\top \hat{\mathbf{H}}^{\ell, m} \hat{\mathbf{X}}^{\ell, m}$ (Eq. 26 in C21). This formulation scales the regularization strength with harmonic mode, applying stronger regularization at higher multipoles (where noise dominates) while preserving large-scale features. For the refinement phase, we use Strategy #4, which determines $\epsilon_{n, \ell}$ from the angular power spectra of the sources $\mathbf{S}$ and noise from the previous iteration, modulated by the hyperparameter c (Eq. 27 in C21).

In all applications presented in this work, the harmonic inversion is restricted to $\ell \leq 200$. The number of sources is set to $n_s = 3, 4$, and 5 for the degraded, evolving, and oscillating beam

Table 1. Parameters adopted in this work.

Beam configuration	Degraded	Evolving	Oscillating
n_s	3	4	5
c	10^{-3}	5×10^{-3}	10^{-2}
$K_{\max}$	0.3	0.8	0.7

configurations, respectively, with increasing n_s , reflecting the rising complexity of the instrumental response. The SDecGMCA algorithm is described in detail in C21. For completeness and reproducibility, the main hyperparameter choices and implementation settings adopted in this work are summarized in Table 1.

SDecGMCA has not previously been applied to HI intensity mapping. In this work, we evaluate its performance in this new context under realistic observational conditions, including frequency-dependent beam distortions and strong foreground contamination.

4. Results

In this section, we apply SDecGMCA to the simulated data (Sect. 2) to assess its performance in reconstructing the cosmological HI 21 cm signal, and compare it with the foreground removal methods described in Sect. 3. We evaluate both the angular and frequency power spectra in order to capture fluctuations of the signal across the sky and along the line of sight. Additional quantitative diagnostics, including the correlation between the true and reconstructed HI maps and the variance of the fractional residual power, $\Delta C_\ell/C_\ell$, as a function of ℓ , are presented in Appendix A.

4.1. Data pre-processing: masking the peaks

Before applying any foreground-removal technique, we introduced a pre-processing step designed to improve foreground identification and, consequently, HI reconstruction. The procedure follows the logic of the polynomial fitting approach. First, we fitted a sixth-order polynomial along each line of sight to model the smooth baseline. Subtracting this fit flattened the spectra, after which we identified outliers via sigma clipping at a threshold of 5σ on the residuals. These outlier pixels, corresponding to bright HI peaks that are further amplified by instrumental noise, were masked and excluded from a second polynomial fit. The refit, performed on the remaining data, provides a less biased baseline estimate. Finally, the masked pixels were replaced by the value of this second-fit baseline, which effectively smooths the foregrounds and facilitates subsequent separation by both BSS and fitting methods. This pre-processing was applied uniformly to all input maps prior to component separation. Further implementation details are given in Appendix B. We find this step essential, as unmitigated peaks in the temperature profiles significantly degrade the performance of all tested methods.

4.2. Results of foreground removal methods

To investigate the impact of instrumental effects, we test three beam models (Sect. 2.2.1): (i) a Gaussian-degraded beam with uniform resolution, (ii) a Gaussian-evolving beam with smooth frequency dependence, and (iii) a Gaussian-oscillating beam with oscillatory frequency dependence. Fig. 1 and 2 show the recovered power spectra for all methods across these beam scenarios.

Note that at high multipoles ($\ell > 200$), the reconstruction with SDecGMCA is limited by the implicit beam inversion that occurs when solving for $\hat{S}$ in the minimization problem (Eq. 21 of C21):

$$\hat{S}_{\ell,m} \leftarrow \left(A^\top \text{diag}(H^\ell)^2 A + \text{diag}(\epsilon_{n_s,\ell}) \right)^{-1} A^\top \text{diag}(\hat{H}_\ell) \hat{X}_{\ell,m}. \quad (23)$$

This inversion is stabilized by a Tikhonov regularization controlled by the hyperparameter $\epsilon_{n_s,\ell}$. While the regularization smooths the foregrounds at scales where the beam response is poor, it also introduces a negative bias in the foreground estimate, which translates into a spurious positive bias in the recovered 21 cm signal. As a result, the regularization itself becomes a source of systematic error. In practice, the deconvolution can only be performed up to the multipole range where the matrix $A^\top \text{diag}(H^\ell)^2 A$ remains invertible without regularization (i.e. $\epsilon_{n_s,\ell} = 0$). This defines a maximum multipole $\ell_{\max}$, which depends on the beam width, beyond which the inversion becomes unstable. In our simulations this corresponds to $\ell_{\max} \simeq 200$, and indeed a prominent spurious peak appears across all beam models for $\ell > 200$. For this reason, we cut in our plots the SDecGMCA-reconstructed angular power spectrum at $\ell = 200$.

4.2.1. Gaussian degraded beam

In the simplest case of a Gaussian degraded beam, all methods perform well, recovering both the angular and frequency power spectra with better than 10% accuracy across a wide range of scales. In this case, we used $n_s = 3$ for all the BSS methods. In the angular domain, the reconstruction remains within 10% accuracy for $3 < \ell < 200$, while in the frequency domain all methods achieve similar accuracy for $k_\nu > 3 \times 10^{-2}$. At larger frequency scales ($k_\nu < 3 \times 10^{-2}$), a small underfitting is observed, reflecting the degeneracy between the long-wavelength foreground fluctuations and the cosmological HI signal, an ambiguity that none of the methods can resolve. Overall, such good performance is expected, as the uniform angular resolution preserves the spectral smoothness of the foregrounds across frequency channels, which facilitates their separation from the cosmological signal. In this case, we used $n_s = 3$ for all the BSS methods.

4.2.2. Gaussian evolving beam

When the beam evolves with frequency, the performance of the methods begins to diverge. In this case, we used $n_s = 4$ for all the BSS methods in order to accommodate the extra spectral features introduced by the beam evolution.

The parametric fitting approach fails entirely, overestimating the angular power spectrum by several orders of magnitude across all scales. This likely occurs because the model, which fits only three free parameters per line of sight, lacks the flexibility to capture the additional spectral distortions induced by the frequency-dependent beam. In contrast, the BSS methods recover the angular power spectrum to within approximately 10% accuracy at $4 \lesssim \ell \lesssim 200$, with the exception of GMCA, which shows a clear excess at intermediate and small multipoles. This degradation is likely related to GMCA's reliance on morphological sparsity, which becomes less effective once the evolving beam disrupts the spatial coherence of the foregrounds across frequencies.

In the frequency domain, all methods except the parametric fit successfully recover the power spectrum down to




Fig. 1. Angular power spectrum of the reconstructed HI 21 cm signal using different methods, each represented by a distinct color. The bottom panel displays the ratio of the reconstructed to the true HI power spectrum for each method. From left to right, the panels correspond to different beam models: Gaussian degraded, Gaussian evolving, and Gaussian oscillating.




Fig. 2. Equivalent to Fig. 1, but showing the frequency power spectrum.

$k_\nu \sim 0.1$, below which deviations increase. At these larger frequency scales (lower k_ν), the beam evolution couples angular and spectral modes, producing scale-dependent distortions that the BSS methods cannot fully disentangle. The polynomial fitting method, by construction, is insensitive to this particular effect. However, it still suffers from the intrinsic degeneracy between the smooth frequency dependence of the foregrounds and the large-scale fluctuations of the cosmological signal.

4.2.3. Gaussian oscillating beam

The most realistic and challenging case is that of the oscillating beam, which further complicates the separation of the cosmological signal from the foregrounds. In this scenario, we increased the number of sources to $n_s = 5$ in order to better capture the enhanced complexity introduced by the beam oscillations. The two fitting approaches (parametric and polynomial) fail entirely, yielding highly biased estimates of both the angular and frequency power spectra. The BSS methods perform comparatively better, although significant limitations remain. For this reason, in the following we focus exclusively on the performance of the BSS methods.

In the angular power spectrum, a strong excess appears at $\ell < 10$, similar to the offset observed in the previous cases but now significantly larger. This excess may arise from artificial residuals introduced during the pre-processing step that masks bright HI sources (Sect. 4.1). In the oscillating beam case, the initial polynomial subtraction may leave residual oscillations that are then amplified. During the subsequent sigma-clipping and re-fitting stage, these residuals can be misidentified as bright HI sources, leading to incorrect replacement of foreground-dominated pixels with smoothed fits. Combined with the frequency-dependent beam evolution, these misattributions can generate large-scale artificial structures in the maps. Since

they are partially interpreted as HI signal, they produce an excess at large angular scales in both the angular and frequency spectra. At intermediate scales ($10 < \ell < 200$), SDecGMCA stands out as the only method capable of recovering the HI signal with an accuracy better than 5%, whereas PCA, FastICA, and GMCA show noticeably larger residuals.

In the frequency domain, all methods display a spurious peak around $k_\nu \sim 0.3$. A similar artifact was also reported by [Matshawule et al. \(2021\)](#) and [Spinelli et al. \(2021\)](#), where it was attributed to beam–foreground coupling. Among the methods tested here, SDecGMCA is the most effective at suppressing this peak, underscoring the importance of accounting for the beam within the separation process. Nevertheless, an excess of power persists at low k_ν across all BSS methods, and in this respect SDecGMCA performs no better than PCA or FastICA.

These results highlight the limitations of both standard BSS techniques and simple fitting approaches when confronted with complex, frequency-dependent instrumental effects. In line with [Spinelli et al. \(2021\)](#) and [Matshawule et al. \(2021\)](#), we find that evolving and oscillating beams severely degrade the performance of traditional methods, producing artifacts in both angular and frequency domains. By contrast, SDecGMCA, which jointly performs source separation and beam deconvolution, substantially mitigates some of these effects, providing the most consistent recovery of the HI signal across all beam scenarios. Nevertheless, the intrinsic instability of beam inversion, already well known in the literature, remains a major limitation. Overall, these findings establish SDecGMCA as a particularly promising method for realistic 21 cm intensity mapping experiments. It is worth noting that in our analysis we applied the method to the full sky without removing or masking any regions, and even under these conditions SDecGMCA shows encouraging performance. As we discuss in Sect. 5, a targeted masking strategy of




Fig. 3. Total MSE of the SDecGMCA reconstruction as a function of the number of sources n_s . The MSE is computed with respect to the ground truth by combining both the frequency and angular power spectra. The curve highlights the trade-off between underfitting at low n_s and overfitting at high n_s , with the minimum occurring around $n_s = 5$ (marked in red), which provides the most balanced reconstruction.

the most contaminated lines of sight can further mitigate residual artifacts in the reconstructed maps.

4.3. Dependence on the number of components

The performance of SDecGMCA, as with other BSS methods, depends not only on the instrumental complexity but also on the internal choices made within the algorithm. One of the most critical hyperparameters is the number of sources (n_s) assumed in the separation model. Since n_s sets the dimensionality of the component space, its value strongly influences the balance between foreground suppression and signal recovery. To assess this dependence, we applied SDecGMCA to the oscillating beam simulations, varying n_s between 2 and 20, and for each run computed the mean squared error (MSE) between the reconstructed and true HI signal in both the frequency and angular power spectra (Fig. 3).

A closer inspection of the reconstructed spectra shows (Fig. 4) that the optimal choice of n_s is somewhat scale-dependent. At large frequency scales (low k_ν), $n_s = 5$ provides the best suppression of excess power, while the localized bump around $k_\nu \sim 0.3$ is more efficiently removed when $n_s = 6$ is used. In the angular domain, differences between models are less pronounced at very large multipoles, where beam effects dominate, but at intermediate scales ($\ell \sim 10^1$ – 10^2) the reconstruction improves noticeably for $n_s = 7$. This behavior reflects the fact that different astrophysical foregrounds and instrumental systematics dominate different scales, requiring varying degrees of model flexibility.

Overall, the most balanced reconstruction across both domains is obtained for $n_s = 5$. As Fig. 3 illustrates, the number of sources plays a role analogous to regularization. Too few sources (e.g. $n_s = 2$ – 3) lack the flexibility needed to capture the complex spectral behavior of the foregrounds, causing residual contamination to leak into the HI signal. Too many (e.g. $n_s \gtrsim 10$) risk overfitting, as the algorithm tends to fragment coherent astrophysical components into multiple modes, reintroducing noise and residuals. This trade-off is clearly visible in the MSE curve, which rises again at high n_s . The importance of carefully choosing the number of sources has also been emphasized in the HI extraction literature, where a range of BSS methods applied to different experimental configurations consistently highlight this

trade-off (e.g., Chapman et al. 2013; Mertens et al. 2018; Carucci et al. 2020; Hothi et al. 2021; Carucci et al. 2024).

In summary, while SDecGMCA is capable of robustly recovering the HI power spectra, its performance remains sensitive to the choice of n_s . A moderate number of sources ($n_s \simeq 5$ – 7) provides the best compromise across scales, although the precise optimum can vary depending on the k_ν or ℓ range of interest. A practical challenge arises when applying the method to real data, where the ground truth is not available and the optimal value of n_s cannot be identified directly. In such cases, indirect diagnostics such as the eigenvalues of the covariance matrix of the signal, the stability of the recovered power spectra, or cross-correlations with external tracers will be needed.

Another important set of hyperparameters in SDecGMCA are those associated with the thresholding procedure. We have not discussed them here in detail, as they are more technical and beyond the scope of this work. In a follow-up study, however, we will show that replacing the current thresholding scheme with a structured convolutional neural network that acts as a learned, wavelet-like multiscale representation (Bonjean et al. 2025) can effectively eliminate the dependence of the method on these hyperparameters.

5. Masking effect

Masking the bright peaks in the frequency profiles of individual lines of sight (Appendix B) helps improve the HI reconstruction. These peaks correspond to bright HI sources and not to extragalactic point sources, which has already been studied in Matshawule et al. (2021). In particular, this masking of the peaks mitigates systematic offsets in the frequency power spectrum, such as the broadband bias across all scales reported in Spinelli et al. (2021) and also observed in our results. While SDecGMCA demonstrates improved performance under these conditions, particularly in the challenging oscillating beam case, there remain regions of the sky where the method fails to accurately reconstruct the underlying HI signal. A likely contributing factor is the inclusion of the Galactic plane and other extremely bright lines of sight in the input maps, which make the separation even more challenging.

In this section, we examine how excluding or attenuating these bright-foreground regions influences the quality of the HI signal reconstruction. Although directly applying a Galactic mask during component separation is currently incompatible with the SDecGMCA framework, we explore two indirect approaches: first, by post-processing the reconstructed HI signal to exclude lines of sight with problematic reconstruction (Sect. 5.1) and second, by artificially dimming the foreground amplitude in and around the Galactic plane to simulate the effect of a Galactic mask (Sect. 5.2).

5.1. Post-masking

To investigate the origin of the spurious peak in the frequency power spectrum and the excess power at low k_ν , we closely examined the reconstructed HI maps from all foreground removal methods. We observed that, for a subset of lines of sight, there is coherent spectral structure as a function of frequency, indicative of leaked foregrounds. To identify which lines of sight contribute to the residuals, we computed the MSE between the reconstructed and true HI signal along each line of sight (MSE_{los}). We then flagged the lines of sight with $|\langle \text{MSE} \rangle - \text{MSE}_{\text{los}}| > \sigma_{\text{MSE}}$, where $\langle \cdot \rangle$ denotes the mean and σ the standard deviation over all lines of sight.




Fig. 4. Reconstructed HI frequency (left) and angular (right) power spectra obtained with SDecGMCA for different choices of the number of sources (n_s , shown in different colors). The black line shows the input HI signal for reference, and the lower panels display the ratio between the reconstruction and the true signal.

The resulting flagged lines of sight are illustrated in Fig. 5, both as spectral profiles (brightness temperature versus frequency) and spatially on the sky. Visually, it is evident that these contaminated lines of sight are located along the Galactic plane, where foreground emission is particularly strong. To evaluate their influence on the reconstructed power spectra, we excluded the flagged regions when computing the frequency and angular power spectra of the recovered HI maps. The results of this analysis are presented in Fig. 6.

Applying this post-processing reduces the amplitude of the peak in the frequency power spectrum further, as shown in the left panel of Fig. 6. Moreover, the HI signal is recovered down to lower k_ν scales, significantly improving the fidelity of the reconstruction. To ensure a fair comparison, we applied the same masking procedure to the outputs of the other methods. However, no significant improvement was observed in these cases. This suggests that their performance degradation is introduced during foreground removal and exceeds what can be mitigated by post-processing.

In addition, we applied the mask described in the previous subsection to the angular power spectrum. Masking the Galactic plane before spectrum estimation is a standard practice in the CMB literature (e.g., Bennett et al. 2003; Planck Collaboration et al. 2016), where regions of bright Galactic emission are excluded to reduce contamination and improve the robustness of the cosmological inference. The resulting angular power spectra, shown in the right panel of Fig. 6, demonstrate the clear benefits of this strategy. Excluding the Galactic plane significantly improves the reconstruction of the HI signal for all BSS methods. Among them, SDecGMCA benefits the most, successfully recovering the HI power spectrum down to smaller multipoles compared to the case without post-processing masking.

5.2. Simulating a Galactic mask

Encouraged by the improvements observed when excluding bright regions from the power spectrum computation, we next explore whether incorporating a masking strategy before apply-

ing the separation methods could yield similar or even stronger benefits. Since direct masking is not currently supported by SDecGMCA, we simulate this scenario by artificially dimming the foregrounds in and around the Galactic plane. This approach mimics an observational strategy that avoids highly contaminated regions of the sky and allows us to evaluate how bright Galactic structures impact the reconstruction quality.

To test this idea in a controlled and continuous way, we applied a smooth Gaussian weighting function in Galactic latitude to the simulated foreground maps. The weights preserve the signal outside $|b| > 15^\circ$ and decrease gradually toward the plane, reaching 0.1 at $b = 0^\circ$. Each frequency map was multiplied by this latitude-dependent weight, after which the maps were smoothed with the telescope beam and instrumental noise was added, following the same procedure as for the unmasked simulations. We then applied all foreground removal methods to these artificially dimmed maps using the Gaussian oscillating beam. The reconstructed spectra are shown in Fig. 7. This mock masking improves the performance of all methods, with the most significant gains seen for PCA, FastICA, and SDecGMCA.

In the frequency domain, the spurious peak persists for PCA and FastICA but is suppressed to below the 2% level for SDecGMCA. This confirms the result already reported by Mashawule et al. (2021), namely that the peak is primarily driven by lines of sight dominated by exceptionally bright foregrounds. Our findings therefore strengthen the case that excluding such regions prior to component separation can substantially improve the recovery of the cosmological signal.

In the angular domain, dimming the Galactic plane improves the performance of PCA, FastICA, and SDecGMCA alike, extending the multipole range over which the HI signal is accurately reconstructed. This result suggests that masking bright regions prior to foreground removal can be beneficial not only for SDecGMCA but also for other BSS methods, particularly for the angular power spectrum.




Fig. 5. Filtering lines of sight prior to computing the power spectrum of the reconstructed HI signal. Left: Brightness temperature of the reconstructed HI signal as a function of frequency for all flagged lines of sight that deviate from the expected spectral flatness. Right: Sky map showing the location of the flagged lines of sight, illustrating the regions excluded by the constructed mask. The percentage of lines of sight that are masked is indicated at the top of the map.




Fig. 6. Impact of post-processing masking on the HI signal reconstruction for the oscillating beam case. Left: Frequency power spectrum of the reconstructed HI signal. The solid black curve shows the true HI signal plus noise, while the dotted red curve represents the SDecGMCA reconstruction before masking. The dashed red curve shows the result after excluding problematic lines of sight, which suppresses the spurious peak and improves recovery at lower k_ν . Other methods are shown in lighter transparent lines for comparison, with minimal improvement observed. Right: Same as the left panel, but for the angular power spectrum of the reconstructed HI signal. The top panel displays the reconstructed power spectra from different methods, while the bottom panel shows the ratio to the true signal.

6. Conclusions

We assessed the performance of SDecGMCA, an extension of DecGMCA to spherical data that jointly performs source separation, for HI 21 cm intensity mapping under increasingly realistic beam conditions. We also compared to widely used and known foreground removal techniques, both model-fitting and BSS methods. Our main findings can be summarized as follows:

- Under the simplest Gaussian degraded beam, all methods recover the HI angular and frequency power spectra with better than 10% accuracy across a broad range of scales, as expected given the preserved spectral smoothness of the foregrounds.
- When the beam evolves smoothly with frequency, the performance of the methods diverges. Parametric fitting fails entirely, GMCA starts degrading at intermediate scales in the angular ($\ell \sim 40$) and frequency ($k_\nu \sim 0.2$) power spectrum, while PCA, FastICA, and SDecGMCA maintain accuracy

over wider ranges in both angular ($4 \lesssim \ell \lesssim 200$) and frequency ($k_\nu > 0.1$) domains.

- In the most realistic and challenging oscillating beam case, the fitting approaches fail catastrophically, and only BSS methods remain viable. Among them, SDecGMCA achieves the most accurate reconstruction, suppressing the spurious peak at $k_\nu \sim 0.3$ and recovering the angular spectrum at intermediate multipoles ($20 < \ell < 200$). Nevertheless, beam inversion remains a fundamental limitation, producing unstable results beyond $\ell \sim 200$.
- The number of sources n_s is a critical hyperparameter for BSS methods. We find that a moderate choice of $n_s \approx 5$ – 7 provides the best balance between underfitting and overfitting, with the precise optimum being scale-dependent. This sensitivity highlights the need for robust, data-driven strategies for setting n_s in real observations, where the ground truth is not available.
- We further explored masking strategies. Post-processing by excluding foreground-contaminated lines of sight substan-




Fig. 7. Frequency (left) and angular (right) power spectra of the reconstructed HI signal obtained with different foreground removal methods, shown in various colors. All methods were applied to simulated maps with artificially dimmed foregrounds. These maps were convolved with a Gaussian oscillating beam prior to signal reconstruction.

tially reduces residual artifacts in SDecGMCA, including the suppression of the spurious peak and improved recovery at low multipoles. Simulating a Galactic mask by dimming the Galactic plane demonstrated that such masking improves the performance of PCA, FastICA, and especially SDecGMCA, extending the multipole range over which the HI signal can be reliably reconstructed.

Taken together, our results highlight both the promise and the limitations of current approaches. Simple fitting methods are inadequate once realistic instrumental effects are introduced, while standard BSS techniques such as PCA, FastICA, and GMCA only partially cope with beam distortions. SDecGMCA stands out as a very promising option, consistently performing across all tested scenarios. However, its dependence on hyperparameters (particularly n_s and the regularization scheme) and the intrinsic instability of beam inversion remain major challenges.

Future work should focus on three directions: (i) incorporating masking strategies directly into the SDecGMCA framework; (ii) developing reliable diagnostics for setting hyperparameters from data; and (iii) refining the thresholding scheme (Bonjean et al., submitted). Even without these refinements, our results show that SDecGMCA already delivers robust reconstructions on full-sky simulations without pre-masking, establishing it as a particularly promising technique for upcoming 21 cm intensity mapping surveys.

Acknowledgements. This work was funded by the TITAN ERA Chair project (contract no. 101086741) within the Horizon Europe Framework Program of the European Commission, the French Agence Nationale de la Recherche (ANR-22-CE31-0014-01 TOSCA) and the French Programme National de Cosmologie et Galaxies (PNCG project CIMES).

References

- Acharya, A., Mertens, F., Ciardi, B., et al. 2023, 21-Cm Signal from the Epoch of Reionization: A Machine Learning Upgrade to Foreground Removal with Gaussian Process Regression
- Alonso, D., Bull, P., Ferreira, P. G., & Santos, M. G. 2015, Monthly Notices of the Royal Astronomical Society, 447, 400
- Alonso, D., Ferreira, P. G., & Santos, M. G. 2014, Monthly Notices of the Royal Astronomical Society, 444, 3183
- Amiri, M., Bandura, K., Chakraborty, A., et al. 2024, ApJ, 963, 23
- Amiri, M., Bandura, K., Chen, T., et al. 2023, ApJ, 947, 16
- Anderson, C. J., Luciw, N. J., Li, Y. C., et al. 2018, MNRAS, 476, 3382
- Ansari, R., Campagne, J. E., Colom, P., et al. 2012, A&A, 540, A129
- Asad, K. M. B., Girard, J. N., de Villiers, M., et al. 2021, MNRAS, 502, 2970
- Battye, R. A., Browne, I. W. A., Dickinson, C., et al. 2013, Monthly Notices of the Royal Astronomical Society, 434, 1239
- Battye, R. A., Davies, R. D., & Weller, J. 2004, MNRAS, 355, 1339
- Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, ApJS, 148, 1
- Bennett, C. L., Smoot, G. F., Hinshaw, G., et al. 1992, ApJ, 396, L7
- Bigot-Sazy, M. A., Dickinson, C., Battye, R. A., et al. 2015, MNRAS, 454, 3240
- Bobin, J., Starck, J.-L., Fadili, J., & Moudden, Y. 2007, IEEE Transactions on Image Processing, 16, 2662
- Bobin, J., Sureau, F., & Starck, J. L. 2016, A&A, 591, A50
- Bobin, J., Sureau, F., Starck, J. L., Rassat, A., & Paykari, P. 2014, A&A, 563, A105
- Bonjean, V., Gkogkou, A., & Starck, J.-L. 2025, A&A, submitted
- Bull, P., Ferreira, P. G., Patel, P., & Santos, M. G. 2015, ApJ, 803, 21
- Carlotti Gerosio, R. & Bobin, J. 2021, Digital Signal Processing, 110, 102946
- Carucci, I. P., Bernal, J. L., Cunningham, S., et al. 2024, arXiv e-prints, arXiv:2412.06750
- Carucci, I. P., Irfan, M. O., & Bobin, J. 2020, Monthly Notices of the Royal Astronomical Society, 499, 304
- Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, Phys. Rev. Lett., 100, 091303
- Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, Monthly Notices of the Royal Astronomical Society, 429, 165
- Chapman, E., Abdalla, F. B., Harker, G., et al. 2012, Monthly Notices of the Royal Astronomical Society, 423, 2518
- Chen, T., Bianco, M., Tolley, E., et al. 2024, Monthly Notices of the Royal Astronomical Society, 532, 2615
- Chen, Z., Cunningham, S., Pourtsidou, A., et al. 2025, ApJS, 279, 19
- Crighton, N. H. M., Murphy, M. T., Prochaska, J. X., et al. 2015, MNRAS, 452, 217
- Cunnington, S., Irfan, M. O., Carucci, I. P., Pourtsidou, A., & Bobin, J. 2021, Monthly Notices of the Royal Astronomical Society, 504, 208
- Cunnington, S., Wolz, L., Bull, P., et al. 2023, MNRAS, 523, 2453
- Cunnington, S., Wolz, L., Pourtsidou, A., & Bacon, D. 2019, MNRAS, 488, 5452
- Delabrouille, J., Betoule, M., Melin, J. B., et al. 2013, A&A, 553, A96
- Draine, B. T. 2011, Physics of the Interstellar and Intergalactic Medium
- Fernández, X., Gim, H. B., van Gorkom, J. H., et al. 2016, ApJ, 824, L1
- Fonseca, J. & Liguori, M. 2021, MNRAS, 504, 267
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, Phys. Rep., 433, 181
- Godinaud, L., Acero, F., Decourchelle, A., & Ballet, J. 2023, A&A, 680, A80
- Gogate, A. R., Verheijen, M. A. W., Deshev, B. Z., et al. 2020, MNRAS, 496, 3531
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
- Hothi, I., Chapman, E., Pritchard, J. R., et al. 2021, MNRAS, 500, 2264
- Hyvarinen, A. 1999, IEEE Transactions on Neural Networks, 10, 626
- Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, MNRAS, 389, 1319
- Jiang, M., Bobin, J., & Starck, J.-L. 2017, arXiv e-prints, arXiv:1703.02650
- Kovetz, E. D., Viero, M. P., Lidz, A., et al. 2017, arXiv e-prints, arXiv:1709.09066
- Li, L.-C., Staveley-Smith, L., & Rhee, J. 2021, Research in Astronomy and Astrophysics, 21, 030
- Liu, A. & Tegmark, M. 2011, Phys. Rev. D, 83, 103006

- Maćkiewicz, A. & Ratajczak, W. 1993, *Computers and Geosciences*, 19, 303
- Makinen, T. L., Lancaster, L., Villaescusa-Navarro, F., et al. 2021, *J. Cosmol. Astropart. Phys.*, 2021, 081
- Marins, A., Abdalla, F. B., Fornazier, K. S. F., et al. 2022, *Foreground Removal and 21 Cm Signal Estimates: Comparing Different Blind Methods for the BINGO Telescope*
- Martin, A. M., Giovanelli, R., Haynes, M. P., & Guzzo, L. 2012, *ApJ*, 750, 38
- Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, *ApJ*, 763, L20
- Matshawule, S. D., Spinelli, M., Santos, M. G., & Ngobese, S. 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 5075
- Mazumder, A., Wolz, L., Chen, Z., et al. 2025, *MNRAS*, 541, 476
- MeerKLASS Collaboration, Chatterjee, S., Fornazier, K., et al. 2025, *MNRAS*, 537, 3632
- Mertens, F. G., Ghosh, A., & Koopmans, L. V. E. 2018, *MNRAS*, 478, 3640
- Miville-Deschênes, M. A., Ysard, N., Lavabre, A., et al. 2008, *A&A*, 490, 1093
- Ni, S., Li, Y., Gao, L.-Y., & Zhang, X. 2022, *ApJ*, 934, 83
- Ni, S., Qiu, Y., Chen, Y., et al. 2024, *Application of Physics-Informed Neural Networks in Removing Telescope Beam Effects*
- Olivari, L. C., Remazeilles, M., & Dickinson, C. 2016, *MNRAS*, 456, 2749
- Paul, S., Santos, M. G., Chen, Z., & Wolz, L. 2023, *arXiv e-prints*, arXiv:2301.11943
- Peterson, J. B., Aleksan, R., Ansari, R., et al. 2009, in *astro2010: The Astronomy and Astrophysics Decadal Survey*, Vol. 2010, 234
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, 594, A25
- Pritchard, J. R. & Loeb, A. 2012, *Rep. Prog. Phys.*, 75, 086901
- Remazeilles, M., Dickinson, C., Banday, A. J., Bigot-Sazy, M. A., & Ghosh, T. 2015, *MNRAS*, 451, 4311
- Santos, M., Bull, P., Alonso, D., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 19
- Santos, M. G., Cooray, A., & Knox, L. 2005, *ApJ*, 625, 575
- Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2015, *Phys. Rev. D*, 91, 083514
- Shi, F., Chang, H., Zhang, L., et al. 2024, *Phys. Rev. D*, 109, 063509
- Soares, P. S., Watkinson, C. A., Cunnington, S., & Pourtsidou, A. 2022, *MNRAS*, 510, 5872
- Spinelli, M., Carucci, I. P., Cunnington, S., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 509, 2048
- Starck, J., Murtagh, F., & Fadili, J. 2016, *Sparse Image and Signal Processing: Wavelets and Related Geometric Multiscale Analysis*, 2nd edn. (Cambridge, UK: Cambridge University Press)
- Starck, J.-L., Fadili, J., & Murtagh, F. 2007, *IEEE Transactions on Image Processing*, 16, 297
- Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *Monthly Notices of the Royal Astronomical Society: Letters*, 434, L46
- Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, *MNRAS*, 469, 2821
- Tramonte, D. & Ma, Y.-Z. 2020, *MNRAS*, 498, 5916
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021, *ApJ*, 915, 71
- Wolz, L., Abdalla, F. B., Blake, C., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 441, 3271
- Wolz, L., Blake, C., Abdalla, F. B., et al. 2017, *Mon. Not. R. Astron. Soc.*, 464, 4938
- Wolz, L., Pourtsidou, A., Masui, K. W., et al. 2022, *MNRAS*, 510, 3495
- Zhang, L., Bunn, E. F., Karakci, A., et al. 2016, *ApJS*, 222, 3
- Zonca, A., Thorne, B., Krachmalnicoff, N., & Borrill, J. 2021, *The Journal of Open Source Software*, 6, 3783

Appendix A: Complementary diagnostics

We present additional quantitative diagnostics for the case of the most challenging beam configuration, the oscillating beam, that further assess the performance of the foreground removal methods in reconstructing the cosmological HI signal. These metrics provide complementary insight beyond the power-spectrum analysis discussed in Sect. 4.

Pearson correlation coefficient per map

As a first additional diagnostic, we compute the Pearson correlation coefficient between the reconstructed and true HI maps at each frequency channel. This coefficient quantifies the degree of linear correlation between the spatial fluctuations of the reconstructed (T^{rec}) and input (T^{true}) HI brightness temperature maps:

$$r(\nu) = \frac{\sum_i^{N_{\text{pix}}} (T_i^{\text{true}}(\nu) - \overline{T^{\text{true}}(\nu)}) (T_i^{\text{rec}}(\nu) - \overline{T^{\text{rec}}(\nu)})}{\sqrt{\sum_i^{N_{\text{pix}}} (T_i^{\text{true}}(\nu) - \overline{T^{\text{true}}(\nu)})^2 \sum_i^{N_{\text{pix}}} (T_i^{\text{rec}}(\nu) - \overline{T^{\text{rec}}(\nu)})^2}}, \quad (\text{A.1})$$

where the summation runs over all sky pixels i , from 1 to N_{pix} . Values of $r(\nu)$ close to unity indicate a faithful recovery of the spatial HI structure at frequency ν , whereas lower values point to residual foreground contamination or signal loss. As shown in Fig. A.1, SDecGMCA consistently achieves the highest correlation across frequencies, demonstrating superior reconstruction fidelity compared to the other methods.




Fig. A.1. Pearson correlation coefficient between the reconstructed and true HI maps as a function of frequency. Values close to unity indicate an accurate recovery of the true HI structure, while lower values reveal residual foreground contamination or signal loss.

Variance of the fractional C_ℓ difference

To complement the map-based correlation analysis, we evaluate the variance of the fractional residual difference in the angular power spectrum,

$$R_\ell = \frac{C_\ell^{\text{true}} - C_\ell^{\text{rec}}}{C_\ell^{\text{true}}}, \quad (\text{A.2})$$

averaged over frequency channels. This quantity characterizes the relative deviation of the reconstructed HI power from the true one at each multipole. We then compute the variance of R_ℓ across all frequency channels as

$$\sigma_{R_\ell}^2 = \langle (R_\ell - \langle R_\ell \rangle)^2 \rangle_{\text{channels}}, \quad (\text{A.3})$$

which provides a quantitative measure of the stability and consistency of each method across scales. Lower values of $\sigma_{R_\ell}^2$ indicate more uniform performance and reduced scale-dependent bias. As illustrated in Fig. A.2, SDecGMCA exhibits the smallest variance at intermediate multipoles ($10 \leq \ell \leq 200$), confirming its robustness and stability in reconstructing the HI signal.




Fig. A.2. Variance of the fractional C_ℓ difference as a function of multipole. Lower values correspond to more uniform and consistent performance.

Appendix B: Masking in frequency space

In this appendix, we describe in detail the pre-processing step used to identify and mask sharp peaks in the temperature profiles, which can bias foreground removal. The procedure begins by fitting a sixth-order polynomial along each line of sight (LOS) in the total observed temperature maps (see top panel of Fig. B.1). Subtracting this polynomial from the original data flattens the spectra, as illustrated in Fig. B.2. To identify and exclude bright peaks, we apply sigma clipping with a threshold of 5σ to the flattened residuals. Pixels flagged as outliers (stars in bottom panel of Fig. B.1) are masked.

A second polynomial fit is then performed on the original total temperature profile, this time excluding the masked pixels. Subtracting this refined baseline yields a more accurate estimate and avoids bias from outliers. Fig. B.2 compares the results of baseline subtraction with and without masking. The blue curve shows the flattened signal when all pixels are included, while the orange curve demonstrates the improved baseline removal obtained after excluding bright outliers.

Finally, in the pre-processing strategy used throughout Sect. 4.1, we replaced the masked pixel values with the corresponding baseline values. This smoothing procedure is illustrated in Fig. B.3, where the black line shows the original total signal and the red line shows the smoothed result. This masking was applied consistently to all input temperature maps prior to running either the fitting-based or BSS-based foreground-removal methods, and it substantially improved the identification of foreground-dominated channels, leading to a more accurate HI reconstruction.




Fig. B.1. Temperature profile for a single line of sight. *Top*: Total observed brightness temperature $T(\nu)$ (black) and the underlying HI signal along the same LOS (red). *Bottom*: Zoom on the frequency range containing the prominent peak (note the logarithmic T scale). Star symbols mark the channels flagged by our sigma-clipping criterion and replaced by the second-pass polynomial baseline (Appendix B). This example corresponds to the line of sight hosting the brightest HI source in the simulation and illustrates the masking procedure.




Fig. B.2. Flattened signal after baseline subtraction. The blue curve shows the result when all pixels are included in the polynomial fitting, while the orange curve shows the improved baseline obtained when bright outlier pixels are masked and excluded from the second fit.




Fig. B.3. Temperature profile along a single LOS. The black line shows the original total signal, while the red line shows the smoothed signal obtained after applying the masking procedure described in Appendix B.