

Hydrogen intensity mapping with MeerKAT: Preserving cosmological signal by optimising contaminant separation[★]

I. P. Carucci^{1,2,★}, J. L. Bernal³, S. Cunnington⁴, M. G. Santos⁵, J. Wang^{6,5}, J. Fonseca^{7,5}, K. Grainge⁴, M. O. Irfan⁸, Y. Li^{9,5}, A. Pourtsidou^{10,11}, M. Spinelli^{12,5}, and L. Wolz⁴

¹ INAF – Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, 34131 Trieste, Italy

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

³ Instituto de Física de Cantabria (IFCA), CSIC-Univ. de Cantabria, Avda. de los Castros s/n, E-39005 Santander, Spain

⁴ Jodrell Bank Centre for Astrophysics, Department of Physics & Astronomy, The University of Manchester, Manchester M13 9PL, UK

⁵ Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, Cape Town 7535, South Africa

⁶ Shanghai Astronomical Observatory, Chinese Academy of Sciences, 80 Nandan Road, Shanghai, 200030, China

⁷ Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto CAUP, Rua das Estrelas, PT4150-762 Porto, Portugal

⁸ Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK

⁹ Liaoning Key Laboratory of Cosmology and Astrophysics, College of Sciences, Northeastern University, Shenyang 110819, China

¹⁰ Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh, EH9 3HJ, UK

¹¹ Higgs Centre for Theoretical Physics, School of Physics and Astronomy, University of Edinburgh, Edinburgh EH9 3FD, UK

¹² Observatoire de la Côte d’Azur, Laboratoire Lagrange, Bd de l’Observatoire, CS 34229, 06304 Nice cedex 4, France

Received 16 December 2024 / Accepted 24 August 2025

ABSTRACT

Removing contaminants is a delicate, yet crucial step in neutral hydrogen (HI) intensity mapping and often considered the technique’s greatest challenge. Here, we address this challenge by analysing HI intensity maps of about 100 deg² at redshift $z \approx 0.4$ collected by the MeerKAT radio telescope, an SKA Observatory (SKAO) precursor, with a combined 10.5-hour observation. Using unsupervised statistical methods, we removed the contaminating foreground emission and systematically tested, step-by-step, some common pre-processing choices to facilitate the cleaning process. We also introduced and tested a novel multiscale approach: the data were redundantly decomposed into subsets referring to different spatial scales (large and small), where the cleaning procedure was performed independently. We confirm the detection of the HI cosmological signal in cross-correlation with an ancillary galactic data set, without the need to correct for signal loss. In the best set-up we achieved, we were able to constrain the HI distribution through the combination of its cosmic abundance (Ω_{HI}) and linear clustering bias (b_{HI}) up to a cross-correlation coefficient (r). We measured $\Omega_{\text{HI}} b_{\text{HI}} r = [0.93 \pm 0.17] \times 10^{-3}$ with a $\approx 6\sigma$ confidence, which is independent of scale cuts at both edges of the probed scale range ($0.04 \lesssim k \lesssim 0.3 h \text{ Mpc}^{-1}$), corroborating its robustness. Our new pipeline has successfully found an optimal compromise in separating contaminants without incurring a catastrophic signal loss. This development instills an added degree of confidence in the outstanding science we can deliver with MeerKAT on the path towards HI intensity mapping surveys with the full SKAO.

Key words. methods: data analysis – methods: statistical – cosmology: observations – large-scale structure of Universe

1. Introduction

Remarkably promising yet challenging to detect, mapping the cosmos with the fluctuations in 21-cm radiation from neutral hydrogen (HI) is set to revolutionise the study of the Universe’s large-scale structure (Bharadwaj et al. 2001; Chang et al. 2008; Loeb & Wyithe 2008). The excellent redshift resolution of intensity maps covering vast areas enables us to characterise the structures’ expansion history and growth, providing constraints on its dark energy, dark matter, and neutrino mass properties (Bull et al. 2015; Villaescusa-Navarro et al. 2017; Carucci 2018; Obuljen et al. 2018; Witzemann et al. 2018; Berti et al. 2022).

However, at the same frequencies of the redshifted 21-cm line, other astrophysical sources, such as our Galaxy, shine with considerably higher intensities (by four to five orders of magnitude), adding gigantic foregrounds to the sought-after sig-

nal (Ansari et al. 2012; Alonso et al. 2014). That huge dynamic range between the cosmological signal and the foregrounds makes any tiny miscalibration or instrumental systematics produce catastrophic leakages, mixing up signals, and rendering their separation rather difficult (Shaw et al. 2015; Carucci et al. 2020; Matshawule et al. 2021; Wang et al. 2022).

Nevertheless, many instruments and collaborations worldwide are actively involved in (or have already completed) HI intensity mapping surveys. Some of them have successfully detected the signal by using secondary tracer information (Pen et al. 2009; Chang et al. 2010; Masui et al. 2013; Switzer et al. 2013; Wolz et al. 2017, 2022; Anderson et al. 2018; Tramonte & Ma 2020; Li et al. 2021a; Cunnington et al. 2023a; Amiri et al. 2023, 2024; MeerKAT Collaboration 2025), keeping pace with more and more stringent upper limits (Ghosh et al. 2011; Chakraborty et al. 2021; Pal et al. 2022; Elahi et al. 2024), and even detecting the signal on its own, albeit at small (non-cosmological) scales (Paul et al. 2023). Finally, other radio telescopes have been (sometimes purposely) planned

[★] On behalf of the MeerKAT Collaboration.

^{★★} Corresponding author: isabella.carucci@inaf.it

to perform radio intensity mapping and either construction or systematic characterisation is ongoing (Xu et al. 2015; Hu et al. 2020; Crichton et al. 2022; Abdalla et al. 2022; Li et al. 2023).

In the observational works cited above, the instruments used are varied (e.g. single-dish radio telescopes, phased array feeds, packed interferometers, and arrays used as a collection of single dishes). Contaminant removal strategies, too, differ substantially. Some teams have opted for foreground ‘avoidance’ and tried to filter the subset of the data with the strongest cosmological signal. Others use what the signal-processing community refers to as ‘blind source separation’ algorithms to try to be as agnostic as possible to the nature of the contaminants. Here, we opt for the latter.

This article focuses on the contaminant removal strategy adopted by the MeerKLASS collaboration, which runs HI intensity map observational campaigns with the MeerKAT telescope (Santos et al. 2016; Wang et al. 2021; MeerKLASS Collaboration 2025). In particular, our starting point is the detection of signal in cross-correlation between MeerKAT radio intensity maps and galaxies from the WiggleZ Dark Energy Survey at redshift $z \approx 0.4$ reported in Cunnington et al. (2023a), which we refer to as CL23 in the remainder of the text.

In theory, astrophysical foregrounds are smooth along frequency, compared to the noisy spectral structure of the HI signal; hence, they are easily separable from the signal, for instance, via a principal component analysis (PCA), holding limited spectral degrees of freedom. In practice, bandpass fluctuations, chromatic beam response, and leakage of polarised foregrounds into the unpolarised signal render the spectrally smooth assumptions invalid. In this work, we address whether the hypotheses behind PCA-like methods are still satisfied with observations and how to pre-process the data to make contaminant removal more efficient.

Building up an efficient cleaning pipeline is challenging. Sometimes, the optimum pre-processing choices are unclear and cannot be validated, given that simulations lag behind the complexity of data. The main idea of this paper is to update the cleaning strategy in CL23, test pre-processing choices, and report the conditions under which we re-detect the signal. We, therefore, use the CL23 detection as a benchmark and test pipeline features directly on observations. In particular, we want to assess and mitigate the main drawbacks that we have so far experienced with a PCA analysis: signal loss (e.g. in CL23, up to 80% of the signal power was reconstructed a posteriori at the largest scale probed) and the ambiguity of the cleaning level choice (which increases the variance linked to the measurement).

We have structured the paper as follows. Section 2 describes the HI intensity mapping observations used in our study. Section 3 presents a mathematical framework for the contaminant subtraction problem and defines our cleaning approach. It also introduces a novel multiscale algorithm, mPCA. Section 4 outlines and justifies the data pre-processing steps. Section 5 details the results of the various contaminant separations performed. Section 6 explains how we compute the power spectra and the theoretical model we assume for their fitting. Section 7 showcases the cross-correlation of the cleaned intensity maps with the WiggleZ galaxies, constraining the clustering amplitudes, with Table 1 and Fig. 16 presenting a summary. Finally, Sect. 8 inspects the cubes cleaned with the new mPCA approach more closely, revealing improved component separation compared to standard PCA. We discuss our results and their implications in Sect. 9 and provide our concluding thoughts in Sect. 10.

2. The data set

The observations we use belong to the MeerKAT single-dish HI intensity mapping pilot survey run in 2019 and have already been described in Wang et al. (2021) and used by Irfan et al. (2022, 2024), Cunnington et al. (2023a), Engelbrecht et al. (2025). Below, we summarise the instrument, observation strategy, and pipeline used to compose the intensity maps.

MeerKAT is a radio telescope in the Upper Karoo desert of South Africa that will eventually become part of the final mid-instrument of the SKAO, SKA-Mid. It comprises 64 dishes of $D_{\text{dish}} = 13.5$ m in diameter that harbour three receivers (UHF, L , and S bands). Our data were obtained with the L -band receiver (900–1670 MHz) using a subset of 250 channels belonging to the range $971 < \nu < 1023$ MHz (corresponding to redshifts $0.39 < z < 0.46$ for the 21-cm emission from HI), which is especially free of radio frequency interference (RFI) (Wang et al. 2021).

MeerKAT single-dish observations follow a fast scanning strategy at constant elevation to minimise the effect of $1/f$ noise and retain a continual ground pickup. Dishes move at 5 arcmin s^{-1} along azimuth, scanning about 10 deg in 100 s. The survey targeted a patch of $\sim 200 \text{ deg}^2$ in the WiggleZ 11 h field ($153^\circ < \text{RA} < 172^\circ$ and $-1^\circ < \text{Dec} < 8^\circ$) (Blake et al. 2010; Drinkwater et al. 2010, 2018), away from the galactic plane (where astrophysical foregrounds are brightest) and overlapping with 4031 WiggleZ galaxies, whose positions we cross-correlate with the intensity maps. Observations occurred during six nights between February and July 2019. After observing a bright-flux calibrator, each complete scan (i.e. observational block) across the sky patch took about 1.5 hours. We repeated it seven times, resulting in 10.5 hours of combined data.

Wang et al. (2021) detailed the data reduction pipeline for the time-ordered data (TOD) and subsequent map-making. It includes three steps for flagging RFI, bandpass and absolute calibration using known point sources (3C 273, 3C 237, and Pictor A), and the calibration of receiver gain fluctuations based on interleaved signal injection from a noise diode to remove the long-term gain fluctuations due to $1/f$ noise (Li et al. 2021b). The calibrated total intensity (in Kelvin) TOD are projected into map space (composed by square pixels of 0.3 deg width) adopting the flat-sky approximation and assuming uncorrelated noise. Finally, the calibrated temperature maps are combined for all scans and dishes at all frequencies by pixel-averaging, yielding the final data cube shown in the left panel of Fig. 1.

Observing in single-dish mode yields a map with resolution inversely proportional to the dish diameter. Specifically, the full-width-half-maximum (FWHM) of the central beam lobe for the channel of frequency, ν , is given by (Matshawule et al. 2021):

$$\theta_{\text{FWHM}}(\nu) = \frac{1.16 c \nu}{D_{\text{dish}}}, \quad (1)$$

with c the speed of light. At our median frequency, it corresponds to $\theta_{\text{FWHM}} = 1.28 \text{ deg}$, namely, the size of 3–4 pixels.

Following CL23, as a result of the map-making and pixel-averaging steps, we estimated the pixel noise variance in the maps, \hat{n} , which is inversely proportional to the number of hits (time) stamps. The inverse of the pixel noise variance is the inverse-variance weight maps, $w_{\text{HI}} = 1/\hat{n}$, that we use later in our analysis.

We used the ‘ABBA’ method (Wang et al. 2021) to remove the sky signal in order to estimate the channel-averaged thermal noise per pixel of the data cube. When considering data within their original footprint, we estimated the noise level to




Fig. 1. *Left panel:* Temperature sky map averaged along the frequency range considered, i.e. $971 < \nu < 1023$ MHz. We highlight the footprint of the WigggleZ galaxies in magenta, the smaller footprint CF where we perform the analysis of this work in cyan, and in yellow the Tukey window function we use for the power spectrum computations (dashed and dotted for the zero and 50% boundaries.). *Right panel:* Normalised histograms of the sky temperature of the data cubes for the original footprint (OF) in solid blue with respect to the cropped CF in dashed orange. The histograms are computed from the average map for each cube, to marginalise the frequency-dependent evolution of the temperature field. The double-peak structure reflects the galactic synchrotron gradient (low versus high RA) present in our sky patch, as shown in the left panel.

be of about 2.8 mK; it decreases to 2 mK when cropping data to a more conservative footprint (Fig. 1, left panel, cyan contour).

3. Statistical framework

We assembled the observed scans so that for each frequency channel, ν , we had a map of the total brightness temperature, X . For each given position on the sky (i.e. each pixel, p), the observed temperature is the sum of the cosmological 21 cm signal from HI (T_{HI}), of the contaminants (T_C), and of uncorrelated noise (T_N):

$$X(\nu, p) = T_C(\nu, p) + T_{\text{HI}}(\nu, p) + T_N(\nu, p). \quad (2)$$

In the separation process, we estimated T_C and subtracted it from the original maps to yield the cleaned maps. We estimated T_C by linearly decomposing it as a sum of N_{fg} foreground components (in the signal-processing literature typically called sources, S) modulated by a frequency-dependent amplitude, A :

$$T_C(\nu, p) = \sum_{i=1}^{N_{\text{fg}}} A_i(\nu) S_i(p). \quad (3)$$

We compressed all the maps in a data cube, \mathbf{X} , with $N_{\text{pix}} \times N_{\text{ch}}$ dimensions, with N_{pix} the number of pixels in each map and N_{ch} the number of maps (frequency channels). By merging Eqs. (2) and (3), we can write \mathbf{X} in its matrix form as

$$\mathbf{X} = \mathbf{A}\mathbf{S} + (\mathbf{H} + \mathbf{N}). \quad (4)$$

We call \mathbf{A} the mixing matrix that regulates the contribution of the N_{fg} components \mathbf{S} in the resulting ‘mixed’ signal. Then, \mathbf{H} is the cosmological signal due to the neutral hydrogen 21-cm emission, to be added to a thermal (Gaussian) noise contribution, \mathbf{N} . It follows that \mathbf{A} has $N_{\text{ch}} \times N_{\text{fg}}$, while \mathbf{S} has $N_{\text{fg}} \times N_{\text{pix}}$ dimensions. The problem of contaminant subtraction reduces to determining the foreground driven $\hat{\mathbf{A}}\hat{\mathbf{S}}$ so that $\mathbf{X}_{\text{clean}} = \mathbf{X} - \hat{\mathbf{A}}\hat{\mathbf{S}}$ can estimate the cosmological HI brightness temperature plus the Gaussian noise contribution as accurately as possible. In practice, we only need to determine $\hat{\mathbf{A}}$, as getting $\mathbf{X}_{\text{clean}}$ is equivalent to projecting $\hat{\mathbf{A}}$ on the data cube \mathbf{X} (Carucci et al. 2020):

$$\mathbf{X}_{\text{clean}} = \mathbf{X} - \hat{\mathbf{A}}(\hat{\mathbf{A}}^T \hat{\mathbf{A}})^{-1} \hat{\mathbf{A}}^T \mathbf{X}; \quad (5)$$

in other words, $\hat{\mathbf{A}}$ acts as a filter on the data set.

Solving Eq. (4) to find the estimate $\hat{\mathbf{A}}$ is an ill-posed inverse problem. We need some extra assumptions to move forward.

Here, we focus on the Singular Value Decomposition (SVD) and the related principal component analysis (PCA). Here, we also introduce the multiscale analyses (dubbed mSVD and mPCA), which treat the spatial scales of the maps independently. Next, we see how $\hat{\mathbf{A}}$ is derived according to the different algorithms used in this work (PCA, SVD, mPCA, and mSVD) and their weighted counterparts (PCAw, SVDw, mPCAw, and mSVDw) when using also the noise information contained in the $w_{\text{HI}}(\nu, p)$ maps in the separation process.

3.1. Principal component analysis and singular value decomposition

Overall, PCA is a dimensionality reduction technique that identifies a meaningful basis (in our case, that of the N_{fg} components) to re-express the data set; the principal components should exhibit the fundamental structure of data. Its use is extensive in scope, spanning multitudinous disciplines (see Jolliffe & Cadima 2016 for a general review). PCA assumes that the $\hat{\mathbf{A}}\hat{\mathbf{S}}$ decomposition maximises decorrelation among the new components, eliminating second-order dependencies. Those new orthogonal components are primarily responsible for the variance in the data, thereby revealing important structural features. It can be demonstrated that the directions corresponding to the largest variances are also associated with higher signal-to-noise ratios (Miranda et al. 2008), which holds true for the contaminants present in our observations.

We can obtain the PCA decomposition as the solution to an eigenvalue-eigenvector problem or equivalent to singular value decomposition (SVD) of the centred data cube since SVD is the generalisation of the eigendecomposition (Shlens 2014). Indeed, when we are not mean-centring the maps, we say we are performing an SVD data decomposition (in the literature, SVD is also referenced as an ‘uncentred PCA’).

In practice, computing the PCA of the data set \mathbf{X} entails (i) subtracting the mean at each frequency to get the centred \mathbf{X}_{ctr} ,

$$X_{\text{ctr}}(\nu_i, p) = X(\nu_i, p) - \frac{1}{N_{\text{pix}}} \sum_{j=1}^{N_{\text{pix}}} X(\nu_i, p_j), \quad (6)$$

and (ii) computing the eigenvectors of the frequency-frequency covariance matrix. Instead, with SVD, we can go directly to step (ii), without carrying out the mean-centring. A priori, we do not know whether PCA or SVD is the best approach in our context; both have been used to characterise and remove

contaminants from HI intensity mapping data sets (e.g. Switzer et al. 2013; Anderson et al. 2018 used SVD, CL23 used PCA). The statistics literature is not definitive either (e.g. see discussion in Miranda et al. 2008), as the best decomposition depends on the specific problem under consideration; it is common to test different algorithms directly on data and check the solutions (see, e.g. Alexandris et al. 2017). Here, we also adopted the latter strategy: whether or not to subtract the mean from maps is one of the pre-processing steps we tested in this work; hence, we refer to the PCA or SVD solutions of the contaminant subtraction procedure throughout the work.

Another common pre-processing choice within the PCA framework in the HI intensity mapping context is the weighted PCA (PCAw). This procedure consists of computing the covariance matrix of the weighted data cube, $w_{\text{HI}}\mathbf{X}$, through the inverse noise of the maps to downplay the influence of the pixels we are less confident about. As first realised by Switzer et al. (2013), it is convenient not to increase the rank of the data matrix to prevent the frequency structure in the weights from altering the data covariance. Hence, we recast $w_{\text{HI}}(\nu, p)$ into its 2D projection $w_{\text{HI}}(p)$, taking the mean along each line of sight. The weights echo the integration time per pixel; hence, their separability in frequency-pixel is genuine¹; with the latter projection, we enforced it to be mathematically true.

In the analysis of this work, we aim to test the ‘weighting’ in combination with the other pre-processing steps and this is why we coupled it to either the PCA or SVD.

Summary of PCA-based pipelines. Schematically, the contaminant subtraction pipeline we followed consists of the following steps.

1. We compute the dot product $\mathbf{R} = \mathbf{D}\mathbf{D}^T$, with

$$\mathbf{D} = \begin{cases} \mathbf{X}, & \text{for SVD,} \\ \mathbf{X}_{\text{ctr}}, & \text{for PCA,} \\ w_{\text{HI}}\mathbf{X}, & \text{for SVDw,} \\ w_{\text{HI}}\mathbf{X}_{\text{ctr}}, & \text{for PCAw.} \end{cases}$$

2. We compute the eigenvectors \mathbf{V} of \mathbf{R} and order them by their eigenvalues, from largest to smallest.
3. We use the first N_{fig} eigenvectors (first N_{fig} columns of \mathbf{V}) to define the mixing matrix $\hat{\mathbf{A}}$; namely, $\hat{\mathbf{A}}$ is a subset of \mathbf{V} :

$$\hat{\mathbf{A}} = \mathbf{V}^{N_{\text{ch}} \times N_{\text{fig}}}.$$

Having defined $\hat{\mathbf{A}}$, the solution $\mathbf{X}_{\text{clean}}$ is obtained through Eq. (5).

3.2. The multiscale approach: Different scales require different treatment

A typical concern when cleaning HI intensity maps with methods based on Eq. (4) is that it becomes necessary to increase the aggressivity of the subtraction (by setting higher N_{fig}) to decrease variance in the large spatial scales of the maps. Meanwhile, on small scales, there are no dramatic changes during this process, resembling an ‘overfitting’ condition (e.g. see Fig. 5 in Wolz et al. 2017). Non-observational studies have also begun to spot (and quantify) this trend as more realistic simulations were involved. For instance, from Fig. 15 in Carucci et al. (2020), we can see that a more aggressive cleaning was employed to recover more signal power at large scales

¹ Especially considering that, in the narrow frequency band we consider, RFI-flagging has been minimal.

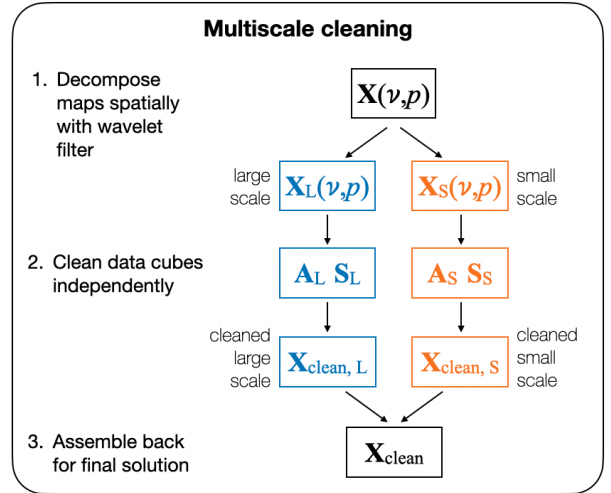


Fig. 2. Flowchart describing the multiscale contaminant subtraction.

at the cost of removing a significant percentage of the signal on the intermediate and small scales. Furthermore, Carucci et al. (2020) demonstrated that the wavelet domain is an optimal, multiscale framework for analysing HI intensity mapping data, enabling a sparse (compact) representation of the astrophysical foregrounds.

Building upon these previous findings, in this work, we want to test the multiscale cleaning with observations. We treated it as an extra pre-processing choice and, thus, we considered a multiscale PCA (mPCA) and the related mSVD and weighted versions, mPCA_w and mSVD_w. As we sketch in Fig. 2, once we decomposed the data cubes in different scales, we solved two independent contaminant subtraction problems, entailing different mixing matrices and allowing for different N_{fig} .

Starlet filtering. The wavelet decomposition offers a straightforward framework for analysing multiscale data. Here, we used the isotropic undecimated wavelet, also known as the starlet transform (Starck et al. 2007), which has proven to be well adapted to an efficient description of astrophysical images (e.g. Flöer et al. 2014; Joseph et al. 2016; Offringa & Smirnov 2017; Peel et al. 2017; Irfan & Bobin 2018; Ajani et al. 2020, and, in particular, Carucci et al. 2020 for the single-dish (low-angular-resolution) HI intensity maps). Starlets have the advantages of (i) representing an exact decomposition; (ii) being non-oscillatory both in real space and in Fourier space, they are efficient at isolating map features represented in either domain; (iii) having compact support in real space, which enables them to prevent systematic effects due to masks and borders.

An intensity map at frequency, ν' , can be decomposed by this transform into a so-called ‘coarse’ version of it, $X_L(\nu', p)$, plus several images of the same size at different resolution wavelet scales, j ,

$$X(p) = X_L(p) + \sum_{j=1}^{j_{\text{max}}} W_j(p),$$

where we dropped the ν' dependency for brevity. Wavelet-filtered maps, W_j , represent the features of the original map at dyadic (powers of two) scales, corresponding to a spatial

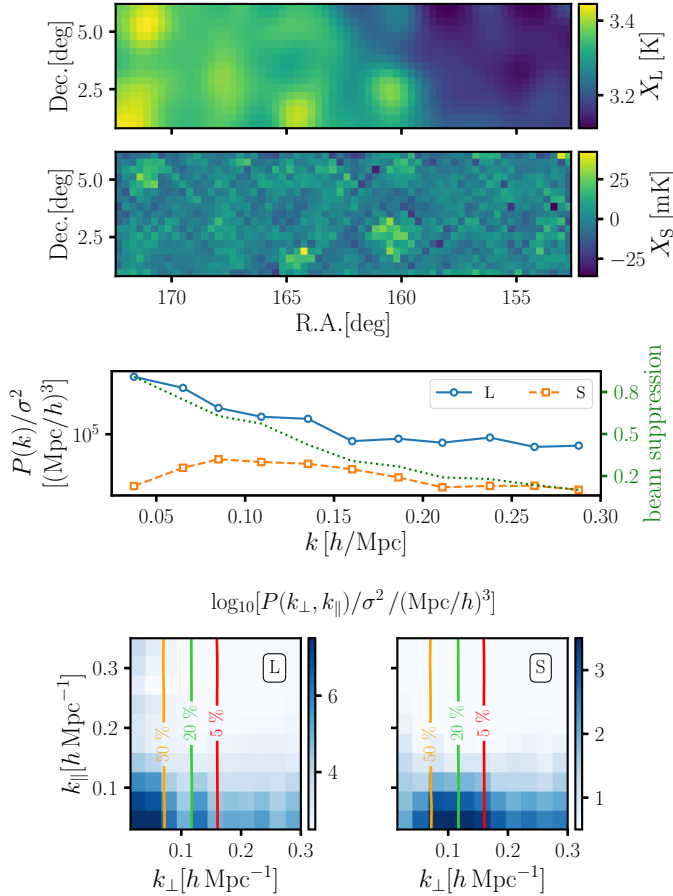


Fig. 3. *First two panels:* Wavelet-filtered large (first panel) and small scale (second) temperature sky map within the conservative footprint (CF) and averaged along all frequencies; i.e. the sum of the two maps above gives precisely the original one (within the cyan contour) in Fig. 1. *Third panel:* Spherically averaged power spectra of the cubes (large scale in solid blue, small in dashed orange) normalised by the variance of each cube. With a green dotted line, we plot the damping term of the telescope beam (refer to the y -axis at the right of the panel). *Bottom panels:* Cylindrical power spectra of the large (left panel) and small-scale cubes (right). Since the wavelet filtering is performed in the angular direction, the difference among the resulting cubes is mostly visible along k_{\perp} (x -axis). The beam suppression also acts in this direction: we plot its expected damping term with iso-contour solid lines corresponding to 50, 20, and 5% suppression, from left to right.

resolution of the size of 2^j pixels, with the largest scale of X_L corresponding to the size of $2^{j_{\max}+1}$ pixels². In our analysis, we set $j_{\max} = 1$ after verifying that using higher wavelet scales does not manifestly improve our results. Hence, applying the starlet decomposition on the observed data cube gives rise to two new cubes,

$$X(v, p) = X_L(v, p) + X_S(v, p), \quad (7)$$

which we refer to as the large-scale and small-scale data sets. We display them in the first two panels of Fig. 3, averaged along frequency, within the smaller ‘conservative’ footprint (later defined in Sect. 4): those maps are the wavelet-filtered equivalents of the cyan-contoured map in the left panel of Fig. 1.

² Due to our narrow frequency band, we overlook the frequency dependence of the maps’ resolution. To achieve an equal distribution of scales, the wavelet filtering should also be frequency-dependent.

Since we went on to later analyse data in three-dimensional Fourier space, we plotted the power spectra of those cubes in the third panel of Fig. 3: the description of ‘large’ and ‘small’ scales holds in k -space, also; although there is no simple, sharp boundary. We compared the behaviour of those three-dimensional spectra with the expected beam suppression with a dotted green line, assuming it is frequency-independent and Gaussian (following Eq. (1) and the modelling later described by Eq. (14)). The values of the damping effect of the beam are on the y -axis on the right of the panel. The large-scale X_L has most of its power for $k \lesssim 0.15 h/\text{Mpc}$, where the telescope beam is suppressing up to 70% of the sky signal, then it flattens. The small-scale X_S lacks power at small k , reaching a peak at $k \approx 0.08 h/\text{Mpc}$, corresponding to $\approx 30\%$ of beam damping. It might be more intuitive to look at the two-dimensional power spectra in the bottom panels of Fig. 3 (left for X_L and right for X_S) since the wavelet filtering is applied in angular space. Most of the power of X_L is at the largest k_{\perp} where the beam damping is below 50%. The X_S power is significant after that 50% bound and slowly decreases.

Summary of the multiscale cleaning. Once we split the data into large-scale and small-scale cubes, we performed two independent contaminant subtractions, as described in Sect. 3.1. These subtractions result in two solutions, $\mathbf{X}_{L,\text{clean}}$ and $\mathbf{X}_{S,\text{clean}}$, which we summed to get the actual cleaned cube of the multiscale analysis, a unique whole-scale solution, as per Eq. (7). We schematically summarise this process in Fig. 2.

The multiscale framework can be coupled to any cleaning method. Here, we focus on the PCA-based algorithms already presented. Hence, when we apply PCA (SVD) to both $X_{L,\text{clean}}$ with N_L components removed and $X_{S,\text{clean}}$ with N_S , we refer to mPCA (mSVD). We can also weight X_L and X_S with the inverse-variance weights, hence resulting in the weighted counterparts, mPCA_w and mSVD_w. We notice from Fig. 3 that the large scale keeps the dimensionality of the cube, so we can apply an SVD to it (i.e. without subtracting the mean temperature value from the maps). The small-scale maps are inherently zero-centred, so using a SVD on them is equivalent to applying a PCA.

3.3. General consequences of the component separation

We expect the cosmological HI signal to be far less correlated in frequency and subdominant in amplitude relative to the foregrounds. These points have two immediate consequences in the separation process.

First of all, (i) the linear decomposition of \mathbf{AS} in Eq. (4) is an inconvenient description of the HI signal and when forcing that decomposition on data, leakage of some of the HI signal into the \mathbf{AS} product is unavoidable. Specifically, if, on the one hand, foregrounds can be restricted to a subset of these modes, on the other, the HI signal spreads across all modes. It, therefore, becomes crucial to perform the separation task with the smallest possible number N_{fg} of components. For instance, [Cunnington et al. \(2021\)](#) report an HI signal loss in a PCA analysis of the order 10% and 60% when N_{fg} is set to 5 and 30, respectively, and finds the latter results almost independent of the complexity of the foregrounds considered in their simulation. Point (i) holds for any separation method based on Eq. (4), although different algorithms may lead to different amounts of signal loss.

Furthermore, (ii) the cosmological signal is inherently coupled to the uncorrelated noise component – such as the thermal part of the instrumental noise – hence, we do not attempt to

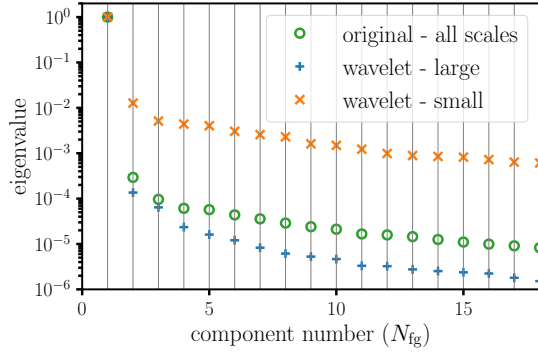


Fig. 4. Normalised eigenvalues of the frequency-frequency covariance matrix of the data cube. Circles refer to the original cube, plus signs to the wavelet-filtered large-scale cube, and ‘x’ crosses for the small-scale cube. The large-scale eigenvalues drop down faster with N_{fg} : the PCA assumption holds better in this case, and a few modes are enough to describe the large-scale data set. These values correspond to mean-centred maps cropped on the conservative footprint, although the eigenvalues behaviour does not change significantly for the other cases.

separate $\mathbf{H} + \mathbf{N}$ at the cleaning stage³. In what follows, with ‘recovered signal’, ‘cleaned data cube’, and similar expressions, we mean the estimate of the cosmological HI signal plus the uncorrelated component of the noise, $\mathbf{X}_{\text{clean}}$.

Considering the number of components N_{fg} . PCA-like methods are dubbed ‘blind’ or ‘unsupervised’: no prior information on the cosmological signal is required. However, we do not know a priori the optimal number of components N_{fg} needed to describe the contaminants and preserve as much cosmological signal as possible.

Typically, as an exploratory analysis, we can plot the eigenvalues of the covariance eigendecomposition in decreasing order, as we do in Fig. 4. The number of plotted values before the last substantial drop in the magnitude of eigenvalues suggests how many components are needed to ‘explain’ the data, that is to say, to account for the largest part of the data variance. Indeed, in the contexts where PCA is successful, the change in magnitude among the eigenvalues is unquestionably visible (in the literature, these plots are called ‘scree plots’). Works based on simulations show evident scree plots; the scree becomes less and less pronounced when additional systematics are modelled on top of astrophysical foregrounds (e.g. Carucci et al. 2020), highlighting the mode-mixing at play. Setting N_{fg} unambiguously in observations is challenging. Methods that automatically find the appropriate sub-space where contaminants live typically need prior information, for instance, on the signal covariance (Shaw et al. 2014; Olivari et al. 2016).

Here, we reverse the argument and claim that if different pre-processing or filtering choices lead to various ‘scree’ eigenvalues behaviour, those that emphasise the magnitude drop are scenarios where the PCA assumptions hold better, resulting into an efficient decomposition. For instance, looking at Fig. 4, the large-scale cube is the data set for which the PCA assumptions hold better (blue plus signs); hence, we can anticipate that on the large scale, PCA will be more efficient – a small N_L will be enough to describe contaminants – compared to the small-scale

counterpart and the PCA without any wavelet filtering. For each cleaning method, we determined the values of N_{fg} (or N_L , N_S) by identifying those that yield the best cross-correlation measurements with the external galactic data set.

4. Pre-processing steps

Here, we discuss the pre-cleaning choices we made to optimise the contaminants-signal separation, highlighting the differences in comparison to the CL23 analysis we have built upon.

PCA-informed footprint. In the left panel of Fig. 1, we show the original footprint, OF, of the data set. With a solid cyan line, we highlight the smaller ‘conservative’ footprint, CF. We have defined the CF by analysing the first four components $\hat{\mathbf{S}}$ found by PCA, projecting $\hat{\mathbf{A}}$ on \mathbf{X} , for the OF analysis; we show these first four components in the first column of Fig. 5. We have drawn iso-contours in normalised intensity in the component maps and found that most of the pixel variance captured by the PCA decomposition was contained in the pixels outside the CF, marked by a solid black boundary. The PCA framework is known to be prone to outlier mishandling⁴, hence the need to suppress their influence either by removing them from the data set, adopting the CF, or by performing a weighted PCAw since the outlier pixels correspond to those with higher noise⁵. A close look at the second column of Fig. 5 suggests that the weighting scheme might not be sufficient and that the complete pixel-flagging choices (last two columns) deliver more reasonable principal components (see later in Sect. 5 a more in-depth discussion on what we mean by a ‘reasonable’ morphology of the sources). We later compare the two options more quantitatively, weighting versus cropping.

We show the temperature histograms of the data cubes in the right panel of Fig. 1. The pixels that have been cut out with the new footprint, after looking at the PCA-derived components, are indeed outliers of the original data temperature distribution.

No extra channel flagging. CL23 performed an extra flagging round on the channels displaying high variance peaks before the cleaning. After channels have been discarded, either we inpaint the missing maps, or we ignore their absence and perform the cleaning on the chopped data cube. CL23 did the latter. However, the PCA decomposition boils down to recovering templates that, modulated in frequency, sum up to the full data cube. When the modulation in frequency is spoiled by the missing channels, the decomposition becomes trickier; in other words, it becomes harder for component analyses of this kind to determine mixing matrices with discontinuities, as shown in Carucci et al. (2020), Soares et al. (2022). Here, we chose to keep all the 250 channels in the analysis. In Appendix A, we checked whether extra flagging could have been beneficial in the set-up of this analysis, ultimately finding it to be detrimental, in agreement with simulation-based works.

Additionally, retaining the full frequency range in the analysis allows us to compensate for the information lost with

³ One possibility is to remove noise after the cleaning by cross-correlating data subsets, since data splits (e.g. in time) contain the same cosmological signal but uncorrelated noise contributions (Switzer et al. 2013; Wolz et al. 2017, 2022; Anderson et al. 2018).

⁴ PCA decomposes data minimising the sum of the squares of deviations from the mean value. The squaring from the outliers dominates the total norm and, therefore, drives the ‘principal’ components of PCA.

⁵ See bottom map of Figure 15 in Wang et al. (2021), that shows the total number of time samples that were combined in each pixel, i.e. inversely proportional to the weights of PCAw and SVDw.

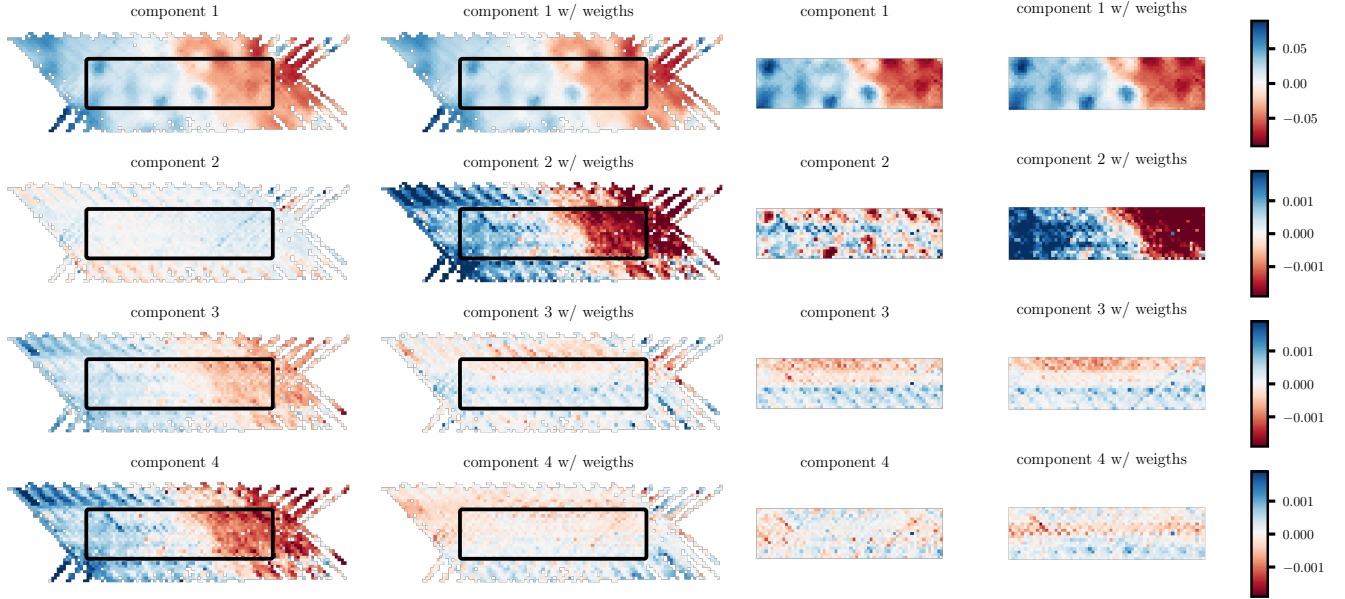


Fig. 5. First four components \hat{S}_{1i} to \hat{S}_{4i} (from top to bottom) in the case where we apply PCA on the original footprint OF (first column) or on the conservative footprint CF (third column) and in the cases where we apply a weighted PCAw (second and fourth columns). All plotted maps are normalised to highlight each component’s spatial features and pixel variance. Component-number (row) wise, we impose the same range of values as per the colour bars on the right. Some sources get saturated, highlighting the higher pixel variance they capture than the unsaturated counterparts (especially for the second component). The black contours in the OF maps in the first two columns mark the CF boundaries.

the CF cropping: we worked with a comparable number of voxels as in CL23, however, with a lower estimated noise (Sect. 2).

No extra smoothing of the maps. Cleaning algorithms based on Eq. (4) perform better when maps share a common resolution. Indeed, in the decomposition, the templates \hat{S} get modulated in frequency through \hat{A} ; however, the mixing matrix \hat{A} cannot accommodate for resolution differences at the map level, as it gives only an overall amplitude adjustment to the \hat{S} components in the frequency direction. Given this intrinsic limitation of the $\hat{A}\hat{S}$ decomposition, accurate beam deconvolution should be ideally performed before the cleaning. CL23 opted for smoothing all maps with a frequency-dependent Gaussian kernel to achieve a common resolution 1.2 times worse than the largest resolution (i.e. with a FWHM that is 1.2 times greater than that of the lowest frequency map, assuming it is Gaussian). CL23 were not the first to opt for the ‘reconvolution’ choice (e.g. Masui et al. 2013; Wolz et al. 2017, 2022; Anderson et al. 2018), and, besides downgrading the maps to a common resolution, the extra smoothing is essentially a low-pass filter that removes the small-scale information in the maps, which one can relate to spurious artifacts. However, the effects of the channel-dependent Gaussian smoothing on the properties of the real-beam-convolved HI field have not been fully understood yet. Spinelli et al. (2022) performed different cleaning methods on simulated data with and without the frequency-dependent smoothing and found the latter practice to be detrimental in all the cases analysed (see also Matshawule et al. 2021 for the effects of a realistic beam in the cleaning process). Here, we opted for the simplest option: leave the maps’ resolutions as they are. Moreover, the beam resolution does not change dramatically within the frequency range we use (1.24–1.31 deg), and the level of uncertainties in our measurement allows us to nevertheless model the theoretical expectation for the three-dimensional

power spectrum with a Gaussian beam relative to the median frequency of the data cube (Sect. 6).

We check in Appendix A what reconvolution could do in terms of the final measurements. We find no improvements and even a detrimental effect in the case of the multiscale cleaning.

5. Results of the component separation

5.1. Performing the cleaning within the conservative footprint

5.1.1. Mixing matrix

We started by applying the PCA cleaning (as per the algorithm described in Sect. 3.1) in the OF and CF. The mixing matrix has all the information for getting to the PCA solution; hence, we started by looking at the derived mixing matrices. In Fig. 6, we show the first two modes (components at the top for the first and bottom panel for the second) of the PCA-like decompositions. Especially for the PCA analysis within the OF (blue solid lines), the second principal mode \hat{A}_{12} (bottom panel) is completely dominated by sharp peaks at specific channels: the PCA decomposition uses this mode to describe the large variance fluctuations in a subset of the maps. In contrast, the behaviour of \hat{A}_{12} gets closer to the expected power-law-like when we take weights into account with PCAw (solid orange line). When analysing the CF (green and red lines), modes 1 and 2 are also spectrally smooth, as the galactic diffuse emissions that should dominate the maps.

While there is no one-to-one relationship between PCA-derived components and actual astrophysical objects, it is noteworthy that the first mode of PCA, when run on the CF, exhibits a frequency behaviour that closely resembles the diffuse synchrotron emission of the Milky Way (Irfan et al. 2022). However, a principal component doing better at isolating the galactic synchrotron does not necessarily indicate that the residuals are cleaner; a first principal component that combines, e.g. free-free and synchrotron emissions, would also reflect the expected




Fig. 6. Normalised first two modes derived by the PCA analysis performed on the original footprint OF (solid blue line) and the conservative footprint CF (solid orange), and their weighted version ‘OFw’ (solid green) and ‘CFw’ (red dashed). For comparison, in the first mode plot (top panel), we show a proxy for the spectral index of the galactic free-free emission (dash-dotted, spectral index of -2.13 , Dickinson et al. 2003) and a range of possible galactic synchrotron spectral indexes (blue shaded area, between -3.2 and -2.6 , Irfan et al. 2022).

dominant contributions in the data. Indeed, all first modes displayed in the top panel of Fig. 6 ‘have something to do’ with the galaxy, which is expected and reassuring.

The SVD mixing matrices (with or without weighting and for the different footprints) are extremely similar to the PCA counterparts; hence, we do not show them here.

5.1.2. Sources

The modes (i.e. the columns of $\hat{\mathbf{A}}$) described above are associated with the corresponding components:

$$\hat{\mathbf{S}} = (\hat{\mathbf{A}}^T \hat{\mathbf{A}})^{-1} \hat{\mathbf{A}}^T \mathbf{X}.$$

We show the first four in Fig. 5, from top to bottom; the first two columns refer to the analysis with the OF, last two columns to the CF, without and with weighting. Components associated with the OF analysis are different than those associated with the CF. Even though the PCA components are not straightforward to interpret as they cannot be exactly paired to physical emissions, they are nevertheless informative – as for their associated frequency behaviour described above. For example, can we see the galactic morphology influence in them? For the OF case, components 1, 3 and 4 display the left-to-right gradient that is linked to the synchrotron at this sky position (Irfan et al. 2022), whereas, for the PCA solution for the CF, the synchrotron gradient is visually present for the first component only. The other three likely mainly capture the morphology of some non-astrophysical systematics, given e.g. the sharp horizontal boundaries⁶ of component 3 and the zigzag structure of component 4. In the CF case, components that look visually galactic are linked to smooth modes (see above and in Fig. 6). In contrast, the situation is more confusing for the OF case: galactic-like components are associated with non-smooth modes, too. Indeed, when working with the OF, we are witnessing more ‘mode-mixing’.

⁶ In Meerclass Collaboration (2025), with further *L*-Band observations, we have noticed that these horizontal stripes are linked to RFI contamination during the downlink of a 50 km away mobile communication tower detected by the beam side-lobes. The development of a method to deal with this effect before the cleaning stage is ongoing.

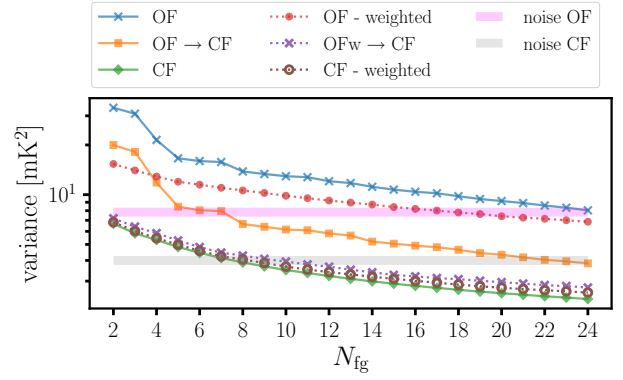


Fig. 7. Variance in the PCA-cleaned cubes as a function of N_{fg} . The solid blue line corresponds to a PCA cleaning applied to the original footprint OF that, if cropped a-posteriori on the smaller CF, corresponds to the variance displayed by the solid orange line. The solid green line corresponds to PCA performed directly on the CF. We compare these results with a weighted PCAw with dotted lines, in the OF case (red), OF-later-cropped (violet) and CF (brown). The horizontal stripes mark the estimated noise level of the OF (pink) and CF (grey) cubes. These results do not change for an SVD analysis.

Again, we do not show SVD results for brevity. The SVD solutions appear somewhat different regarding the $\hat{\mathbf{A}}\hat{\mathbf{S}}$ decomposition; however, these differences are not appreciable enough to make claims at this stage of whether the SVD analysis is more or less optimal than PCA (e.g. the pre-processing choice of using the CF rather than the OF has a much more significant impact). We let the cross-correlation measurements quantify the PCA-SVD distinctness.

5.1.3. Variance

We expect the variance to decrease as a function of N_{fg} . Furthermore, to minimise cosmological signal loss that increases with N_{fg} , we would aim to reach the expected noise level of the data with the smallest N_{fg} possible. We recall here that noise dominates the cosmological signal in this data set. Hence, even though this exercise cannot tell us how much residual contamination is still present in the cubes, we can compare the PCA-based pipelines as the faster (i.e. lower N_{fg}), they reach the noise level, the more optimal. We do so with the solid lines in Fig. 7. The blue line, corresponding to a PCA analysis on the original footprint OF, reaches the OF noise level (pink stripe) at $N_{\text{fg}} = 23-24$; whereas the green line (PCA on the CF) reaches the CF noise level (grey) at $N_{\text{fg}} = 7-8$. Performing the PCA decomposition on the CF leads to a more efficient signal component separation.

By working on the CF, we are suppressing the high variance pixels at the beginning of our analysis, namely, before the cleaning; thus, we might want to consider whether: (i) the previous result is due to the missing outlier pixels rather than a more efficient cleaning and (ii) performing a weighted PCAw on the full OF could allow us to suppress the influence of the high-variance pixels in the signal separation, without cropping the maps. We tested both hypotheses and checked whether they lead to a cleaning efficiency comparable to (or even better than) the PCA applied on the CF.

Regarding hypothesis (i), we cropped the cleaned cubes obtained with the PCA on the OF to the CF, by trimming the maps after the cleaning. The results are shown in orange in Fig. 7. We note that we are not improving the OF results, as we keep reaching the noise (in grey for the CF) around $N_{\text{fg}} = 21-22$.

The extra channel flagging is a condition for PCA to deliver a better decomposition. The filtering solution differs from the OF case for the retained pixels, as previously checked by looking at the modes and components of the decomposition.

To check hypothesis (ii), we employed PCAw. The results are shown by the red dotted line in Fig. 7: we improve upon the plain PCA solution (blue line) by being more effective and reaching the noise floor at $N_{\text{fg}} = 15\text{--}16$, but not as efficiently as in the CF case (green line). Instead, performing a weighted PCAw for the CF case (dotted violet line) does not improve upon the plain PCA. The weight map inside the CF is indeed more uniform, making it less worthwhile to apply the weighting.

We have demonstrated that if data do not have an approximately uniform weight across the cube, then a weighted decomposition is worthwhile in the PCA framework, although it is not as crucial as removing the outlier pixels in the temperature distribution altogether.

Summarising this section, we have found that a PCA analysis on the small CF is more effective than the OF and than its weighted version. Similar conclusions apply to SVD. We continued working solely with the CF in the remainder of this work.

5.2. Cleaning spatial scales independently

In this section, we take a closer look at mixing matrices, sources, and variance of the multiscale cleaning solutions. We focus mostly on mPCA, as we find no appreciable differences with respect to mSVD.

5.2.1. Mixing matrix

Overall, mPCA corresponds effectively to two PCA analyses, PCA-L and PCA-S, performed independently on the maps' large and small spatial scales and acting on them separately. We started by looking at the two mixing matrices, $\hat{\mathbf{A}}_{\text{L}}$ and $\hat{\mathbf{A}}_{\text{S}}$, and see which modes they will project out from the data. We plot the first four modes of both matrices in Fig. 8, solid blue and dotted green for $\hat{\mathbf{A}}_{\text{L}}$ and $\hat{\mathbf{A}}_{\text{S}}$ respectively, and compare them with the plain PCA solution too (empty orange circles). Each panel from top to bottom refers to a mode. The first two modes are almost identical for PCA and PCA-L; instead, the last two differ mainly in the position of the peaks. Since these modes are projected on the entire data cube for PCA and only on the coarse scale of the maps for PCA-L, this finding implies that, in the case of the plain PCA, the large-scale information drives the decomposition, imposing its solution (filtering) also on the small scale of the maps. Indeed, the first two modes of the small-scale driven PCA-S follow the same trend but are much more fluctuating: they are accommodating for channel amplitude fluctuations of these first modes occurring just on the small scale.

5.2.2. Sources

The modes discussed above correspond to the components $\hat{\mathbf{S}}_{\text{L}}$ and $\hat{\mathbf{S}}_{\text{S}}$ that we show in Fig. 9, where the left column is for the large-scale analysis and the rest for the small-scale. We can compare these maps with the whole-scale $\hat{\mathbf{S}}$ components shown in the third column of Fig. 5. The $\hat{\mathbf{S}}_{\text{L}}$ maps indeed reflect the large scale of \mathbf{S} , namely, they are coarse-grained versions of the \mathbf{S} maps. Again, this confirms the hypotheses that (i) the large-scale information drives the plain PCA analysis when working with the whole-scale original cube; (ii) the smooth astrophysical emissions acting as foregrounds correspond to morphologi-




Fig. 8. Normalised first four modes (from the top to the bottom panel) derived by the PCA analysis (empty orange circles) and the multiscale mPCA (lines). With mPCA, we derive two mixing matrices, $\hat{\mathbf{A}}_{\text{L}}$ for the large and $\hat{\mathbf{A}}_{\text{S}}$ for the small scales, and we show their first modes with solid blue and dotted green lines respectively. The first two modes of all three mixing matrices look correlated, as the main astrophysical foregrounds are driving the first modes at large and small scales alike. Nevertheless, the extra fluctuations of the $\hat{\mathbf{A}}_{\text{S}}$ first mode is a hint that the decomposition at the small-scale is characterising extra contributions more relevant at this spatial scale than at the large one.

cally wide, large-scale regions, compared to the sky area and the intrinsic spatial resolution of the observed maps.

It is harder to interpret the $\hat{\mathbf{S}}_{\text{S}}$. They are indeed descriptions of the small scale decomposition, breaking down the ‘stripes’ and ‘zigzag’ seen in $\hat{\mathbf{S}}$; except maybe the first $\hat{\mathbf{S}}_{\text{S}}$ component, where we do see hints of the point sources present in our footprint (Wang et al. 2021). It is indeed this first $\hat{\mathbf{S}}_{\text{S}}$ that corresponds to a mode with a similar trend with what is seen in the mixing matrices $\hat{\mathbf{A}}$ and $\hat{\mathbf{A}}_{\text{L}}$: we expect it to carry some information of the astrophysical foregrounds. To a lesser extent, this is true also for the second $\hat{\mathbf{S}}_{\text{S}}$ component: it is associated with a smooth mode, close to the second mode of $\hat{\mathbf{A}}$ and $\hat{\mathbf{A}}_{\text{L}}$; yet in this case we cannot recognise an astrophysically motivated morphology. Probably, the latter is because these astrophysical emissions have leaked in a systematics-driven component. For concision, we do not show the modes 5–8 associated with those last four $\hat{\mathbf{S}}_{\text{S}}$ components: they are spiky and oscillating around zero (modes with short correlation in frequency), so it is hard to relate them to any motivated signal component or systematic contaminant. We might ask whether these small-scale fluctuating components are actual cosmological signal. Even though we always incur some signal loss, these are still the first eight principal components, those of highest variance, and typically smoother than components that would be next in the decomposition. The PCA-S is no more and no less harmful in terms of signal loss than PCA-L. It is

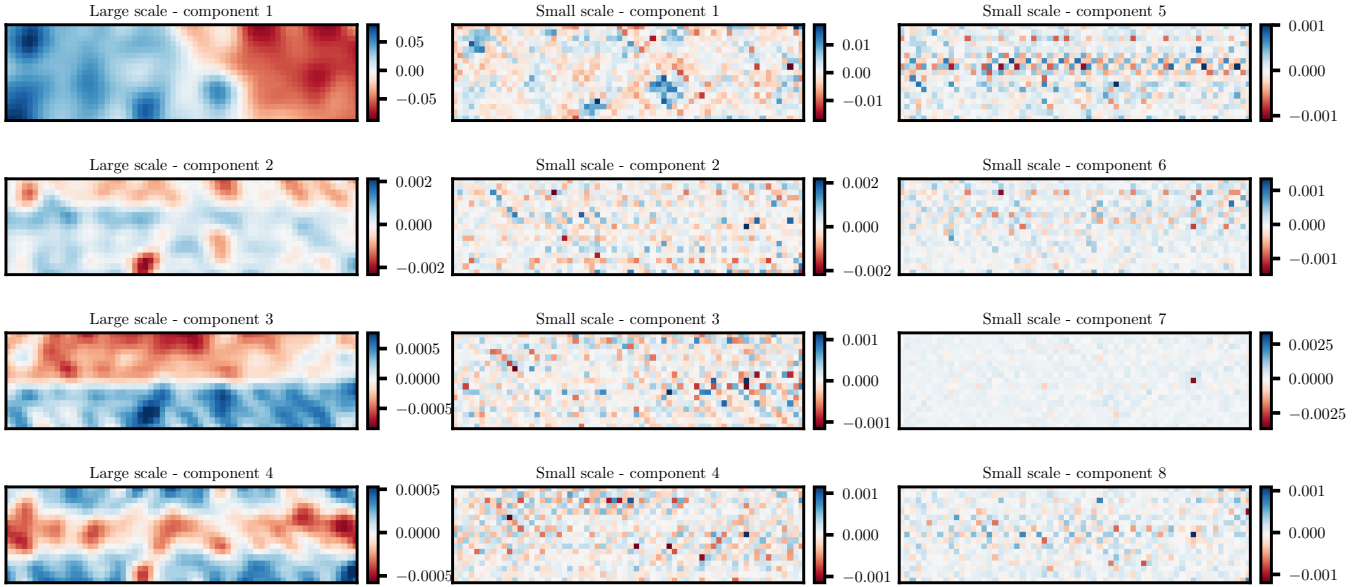


Fig. 9. First four normalised components \hat{S}_L in the case we apply PCA on the large-scale only (left column) and the first eight \hat{S}_S on a small scale. Large-scale components (left column) reflect the coarse structure of the whole-scale PCA decomposition (third column of Fig. 5), whereas, the small-scale components, except the first one to a little extent, do not show any recognisable astrophysical ‘morphology’ in them: they characterise other, small-scale, features, equally crucial in the PCA decomposition (those are the first eight principal components).




Fig. 10. Variance in the mPCA cleaned cubes as a function of the number of components removed when setting $N_L = N_S (=N_{fg})$, as in the x -axis). The legend is the same as in Fig. 7. These results do not change for an mSVD analysis.

just acting on smaller scales, as we explicitly check in Sect. 7.6, where we test on simulations the filtering effect of mPCA.

5.2.3. Variance

We look at the variance of the cleaned mPCA cubes as a function of the number of components removed. Although we decided to work with the CF footprint only, in Fig. 10 we show the variance of the OF cubes cleaned with mPCA, to check the effects of cropping and of weighting for mPCA (and mSVD) too. As in Fig. 7, the solid line in blue (green) corresponds to an mPCA cleaning applied to the OF (CF), to be compared with the weighted counterpart in dotted red (brown) and to the effect of cropping the OF cleaned results to the CF (solid orange of mPCA and dotted violet for weighted-mPCA); pink (grey) bands are the expected noise floor for the OF (CF) observed cubes. On the x -axis we show the number of modes removed, fixing $N_L = N_S$. The con-




Fig. 11. Variance in the mPCA cleaned cubes varying the number of components removed. With dotted lines and filled symbols, we refer to cleaning by fixing the number of the large-scale components removed, N_L , given by the colour of the line (inset colour scale), and with solid lines and empty symbols, the other way around: N_S is kept fixed (colour) and N_L is varying (x -axis). With empty magenta symbols we highlight the cubes cleaned with $N_L = N_S$, for a more direct comparison with the PCA results, shown by the red squares and solid line. We show the estimated uncorrelated noise level with a grey band.

clusions that we can draw are similar (but stronger) to those we drew from Fig. 7: (i) weighting alone is not sufficient for dealing with the high-variance pixels, and (ii) the cropping is indeed beneficial for the mPCA solution to be more descriptive hence efficient for our data cube. We do not show any other weighted mPCAw (or mSVD, mSVDw) results, as we do not find appreciable differences with respect to mPCA.

It is informative to compare Figs. 7 and 10: our first direct comparison of the multiscale cleaning versus the standard PCA. Even when working within the OF, the mPCA residual cube reaches the noise floor (in pink) with only $N_{fg} = 12-13$

components removed. These results do not take full advantage of the multiscale method that entails two independent subtractions, since we are imposing $N_L = N_S$.

We look at the effect of varying N_L and N_S independently in Fig. 11, for the CF only. With the dotted lines, we let N_S vary and keep N_L fixed, given by the colour of the line (inset colour scale); it is the other way around for the solid lines (N_L varying, N_S fixed). Increasing N_L is the primary driver of a variance drop for $N_L \leq 4$; after four modes removed on the large scale, the variance decrease slows down (keeping N_S fixed). Conversely, the contribution of the increase of N_S to reduce the variance is more constant, also beyond the range of mode numbers explored in Fig. 11. With the empty magenta symbols, we highlight the cases where the cleaning has been performed with $N_L = N_S$ for a quicker comparison with the PCA results (solid red line). We demonstrate once again that applying PCA separately on the spatial scales (i.e. mPCA) is not equivalent to applying it on the whole-scale original cube, even when the number of components removed is the same.

The results shown so far have been primarily qualitative, looking at the properties of the PCA and mPCA decompositions. In Sect. 9, we discuss how demonstrative these tests are and what we have learnt from them. In the next section, we present more quantitative results and evaluate the cleaned cubes by measuring their cross-correlation signal with the WiggleZ galaxies.

6. Cross-power spectrum: Estimation and theoretical modelling

We analysed the cosmological signal present in the maps through its cross-power spectrum with the position of the 4031 WiggleZ galaxies, $P_{\text{HI,g}}$. For the estimator and theoretical modelling, we closely followed CL23.

6.1. Cross-power spectrum estimation

The MeerKAT data cube is in celestial coordinates and receiver channel frequency (RA, Dec, ν). To perform fast Fourier transform (FFT) calculations on data, we re-grid the cube onto regular three-dimensional Cartesian coordinates \mathbf{x} with lengths l_{x_1} , l_{x_2} , l_{x_3} in h^{-1} Mpc, assuming a *Planck18* cosmology (Planck Collaboration I 2020) and voxel volume, V_{cell} .

The Fourier transforms of the HI temperature maps δT_{HI} and the galaxy number field n_g are defined as

$$\tilde{F}_{\text{HI}}(\mathbf{k}) = \sum_{\mathbf{x}} \delta T_{\text{HI}}(\mathbf{x}) w_{\text{HI}}(\mathbf{x}) \exp(i\mathbf{k} \cdot \mathbf{x}), \quad (8)$$

$$\tilde{F}_g(\mathbf{k}) = \sum_{\mathbf{x}} n_g(\mathbf{x}) w_g(\mathbf{x}) \exp(i\mathbf{k} \cdot \mathbf{x}) - N_g \bar{W}_g(\mathbf{k}), \quad (9)$$

where $N_g = \sum n_g$ is the total number of galaxies, \bar{W}_g is the weighted Fourier transform of the normalised selection function, W_g , which takes care of the incompleteness of the galaxy survey, as designed by the WiggleZ collaboration (Blake et al. 2010), and the weights $w_{\text{HI}}(\mathbf{x})$ and $w_g(\mathbf{x})$ are the inverse variance map introduced in Sect. 2 and the optimal weighting of Feldman et al. (1994), for the HI and galaxy fields respectively.

Additionally, to limit the ringing in the spectra calculations (i.e. spurious correlations between adjacent k -bins), we apply tapering functions that smoothly suppress the cleaned data to zero at the edges. CL23 apodised the cube in the frequency direction using a Blackman window (Blackman & Tukey 1958); in the angular direction, the weight maps $w_{\text{HI}}(\mathbf{x})$ fall off at the

edges due to scan coverage and effectively act as tapering, hence no further tapering was needed. In this analysis, we keep the Blackman choice in the line-of-sight direction, but since we use a smaller footprint – the CF – the weights within the CF no longer act as tapering. Instead, we adopt a Tukey window function⁷ to maximise the number of pixels whose intensity information is retained in the power spectrum computation (Harris 1978); we display it in the left panel of Fig. 1 with the yellow curves—dotted and dashed for when the window reaches 0.5 and 0, respectively.

Finally, we define the cross-power spectrum estimator as

$$\hat{P}_{\text{HI,g}}(\mathbf{k}) = \frac{V_{\text{cell}}}{\sum_{\mathbf{x}} w_{\text{HI}}(\mathbf{x}) w_g(\mathbf{x}) W_g(\mathbf{x})} \text{Re} \left\{ \tilde{F}_{\text{HI}}(\mathbf{k}) \cdot \tilde{F}_g^*(\mathbf{k}) \right\} \frac{1}{N_g}. \quad (10)$$

All spectra are averaged into bandpowers $|\mathbf{k}| \equiv k$ yielding the final spherically averaged power spectra measurements that we compare against theory.

Our cleaned intensity maps are highly noise-dominated (see Sect. 2, i.e. the noise level is of about 2–3 mK), hence we use Gaussian statistics to analytically estimate the uncertainties of the cross-spectrum measurements:

$$\hat{\sigma}_{\text{HI,g}}(k) = \frac{1}{\sqrt{2N_m(k)}} \sqrt{\hat{P}_{\text{HI,g}}^2(k) + \hat{P}_{\text{HI}}(k) \left(\hat{P}_g(k) + \frac{1}{\bar{n}_g} \right)}, \quad (11)$$

where N_m is the number of modes in each k -bin and $\bar{n}_g = N_g / (l_{x_1} \times l_{x_2} \times l_{x_3})$ is the number density of galaxies. The auto-spectra of the fields, \hat{P}_{HI} and \hat{P}_g , are computed in the same fashion of Eq. (10), namely,

$$\hat{P}_{\text{HI}}(\mathbf{k}) = \frac{V_{\text{cell}}}{\sum_{\mathbf{x}} w_{\text{HI}}^2(\mathbf{x})} |\tilde{F}_{\text{HI}}(\mathbf{k})|^2, \quad (12)$$

$$\hat{P}_g(k) = \frac{V_{\text{cell}}}{\sum_{\mathbf{x}} w_g^2(\mathbf{x}) W_g^2(\mathbf{x})} \left[|\tilde{F}_g(\mathbf{k})|^2 - N_g \sum_{\mathbf{x}} w_g^2(\mathbf{x}) W_g(\mathbf{x}) \right] \frac{1}{N_g^2}, \quad (13)$$

where we included a shot noise term in the galaxy auto-spectrum, whereas \hat{P}_{HI} includes (and is dominated by) thermal noise. CL23 compared these analytical error estimations to the ones calculated from cross-correlating the MeerKAT data with the random WiggleZ catalogues used to derive the selection function, finding good agreement. In Sect. 7, we further test the Gaussian error assumption by means of a reduced χ^2 analysis and comparing them directly with the data-driven jackknife errors, finding also good agreement.

6.2. Theoretical modelling

Our fiducial cosmological model is *Planck18*. We assume it to compute the linear matter power spectrum P_m , derived through CAMB (Lewis et al. 2000), and all other cosmological quantities (see below). We work at the effective redshift $z_{\text{eff}} = 0.42$, corresponding to the median frequency of the data cube.

We adopt the following cross-power spectrum model:

$$P_{\text{HI,g}}(\mathbf{k}) = \bar{T}_{\text{HI}} b_{\text{HI}} b_g r (1 + f\mu^2)^2 P_m(k) \times \exp \left[\frac{-(1 - \mu^2) k^2 R_{\text{beam}}^2}{2} \right]. \quad (14)$$

⁷ The Blackman window function is more optimal and designed to have close to the minimal leakage possible. However, in the angular direction, we cannot afford it with the present number of pixels; we will improve on this with future observations with a larger area coverage.




Fig. 12. Cross-power spectra between WiggleZ galaxies and MeerKAT HI intensity maps at redshift $z \approx 0.4$ for the PCA and PCAw cleaned cubes in the top and bottom panel, respectively. Maps are cleaned by removing different numbers of PCA (and PCAw) modes, from dark, $N_{\text{fg}} = 4$, to light colour lines, $N_{\text{fg}} = 9$; 1σ error bars are shown with the corresponding shaded areas. The red dashed line is the fitted model from CL23 (same data, different cleaning strategy) that instead includes a \mathcal{T} -correction for the signal loss; as a rule-of-thumb, the uncorrected CL23 best-fit would show an 80% lower amplitude at the largest scale probed (e.g. below the 1σ boundary of the $N_{\text{fg}} = 4$ and 5 PCAw measurements for $k \lesssim 0.12 h/\text{Mpc}$ in the bottom panel). We plot the product $kP(k)$ to accentuate the small scales that would otherwise be flattened by the much higher low- k measurements.

\bar{T}_{HI} is the mean HI temperature of the field that, at fixed cosmology, depends on the overall neutral hydrogen cosmic abundance Ω_{HI} (Furlanetto et al. 2006):

$$\bar{T}_{\text{HI}}(z) = 180 \Omega_{\text{HI}}(z) h \frac{(1+z)^2}{\sqrt{\Omega_{\text{m}}(1+z)^3 + \Omega_{\Lambda}}} \text{ mK}, \quad (15)$$

where Ω_{m} and Ω_{Λ} are the density fractions for matter and the cosmological constant, respectively. b_{HI} and b_{g} are the HI and galaxy large-scale linear biases and r is their cross-correlation coefficient. We account for linear matter redshift-space distortions (RSD) with the $(1 + f\mu^2)^2$ factor (Kaiser 1987), where f is the growth rate of structure and μ is the cosine of the angle from the line-of-sight. The Gaussian damping in the last term approximates the smoothing effect of the telescope beam; in particular, we compute the corresponding standard deviation (Matshawule et al. 2021)

$$\sigma_{\text{beam}}(\nu) = \frac{\theta_{\text{FWHM}}(\nu)}{2\sqrt{2\ln 2}}, \quad (16)$$

through the angular resolution FWHM parameter θ_{FWHM} of Eq. (1), and multiply it by the comoving angular diameter distance to the effective redshift–median frequency–of the observed data set, yielding $R_{\text{beam}} = 10.8 \text{ Mpc } h^{-1}$. Given the current noise in the data and uncertainties of the measurements, we find the above modelling of beam smearing to be sufficient for this analysis.

For a direct comparison, we sample the power spectrum template in Eq. (14) onto the same Fourier-space grid as the data, and




Fig. 13. Cross-power spectra between WiggleZ galaxies and MeerKAT HI intensity maps cleaned by mPCA at redshift $z \approx 0.4$. We keep the number of modes removed at small scale fixed at $N_{\text{s}} = 4$ (top panel) and $N_{\text{s}} = 8$ (bottom), and we remove different numbers of large-scale modes (from dark, $N_{\text{L}} = 4$, to light colour lines, $N_{\text{L}} = 9$), with 1σ error bars (corresponding shaded areas). The red dashed line is the fitted model from CL23 (same data, different cleaning strategy) that instead includes a \mathcal{T} correction for the signal loss; as a rule-of-thumb, the uncorrected CL23 best-fit would show an 80% lower amplitude at the largest scale probed (e.g. below the 1σ boundary of the $N_{\text{L}} = 4$, measurements for $k \lesssim 0.14 h/\text{Mpc}$). We plot the product $kP(k)$ to accentuate the small scales that would otherwise be flattened by the much higher low- k measurements.

convolve it with the same window functions. Together with fixing cosmology, we set the galaxy bias $b_{\text{g}} = \sqrt{0.83}$ (Blake et al. 2011). The only free parameter left to be fitted against data is the combination $\Omega_{\text{HI}} b_{\text{HI}} r$, namely, the amplitude of the cross-power spectrum of Eq. (14).

For reference, CL23 measured $\Omega_{\text{HI}} b_{\text{HI}} r = (0.86 \pm 0.10) \times 10^{-3}$, derived using a transfer function (\mathcal{T}) approach to compensate for the signal loss, namely, enhancing the measurements by about 80% for $k < 0.15 h/\text{Mpc}$ down to 20% at the smallest scale considered. We did not perform the \mathcal{T} -correction on the spectra measured in this analysis, for reasons that will become clear in later sections, yet delivering a cross-spectra in agreement (and even with higher amplitude) with CL23.

7. Capitalising on our benchmark: Measuring the correlation with the WiggleZ galaxies

We assessed the new cleaning pipeline choices and methods described in this work, confronting the resulting cleaned cubes with the cross-correlation measurement already reported in CL23. To focus on the cleaning strategy's influence on the results, here we use the same tools presented in CL23 for computing and modelling $P_{\text{HI,g}}(k)$, fitting its amplitude, and evaluating the goodness of the fit. The only exception to the previous statement is related to using a 2D window function before the Fourier transform that replaces the effective tapering due to the weight map of the OF (see Sect. 6).

In many figures in this section, we report the CL23 result as a reference. In particular, we show the fitted model of Eq. (14)

with the amplitude derived by CL23 – not the actual $P_{\text{Hr,g}}(k)$ data points – for clarity and to overcome the fact that in CL23 maps had been smoothed, effectively reducing their native resolutions that instead we keep here. However, the reader has to keep in mind that CL23 derived the final power spectrum with the transfer function, \mathcal{T} , procedure, whereas we stress that none of the results of the present work have a \mathcal{T} -correction applied. It might seem an unfair comparison, but it will make our point: we have improved the cleaning procedure so that the cosmological signal loss is more under control.

7.1. The cross-power spectrum

We start with the PCA-like methods' cleaned cubes. We show the cross- $P(k)$ in Fig. 12 in the top for a simple PCA analysis, in the bottom for its weighted version PCAw. Both SVD results, weighted and not, are in between those presented; that is to say: PCA and PCAw bracket the range of the PCA-like analyses. In each panel, we show the cross- $P(k)$ for different cubes cleaned with different numbers N_{fg} of removed components, colour coded from dark ($N_{\text{fg}} = 4$) to light ($N_{\text{fg}} = 9$). The lines and empty symbols are the measurements, and the corresponding shaded areas are their 1σ uncertainties. Besides the encouraging agreement of some of these results with CL23 (red dashed line), the overall behaviour of these power spectra is promising (especially for PCAw) as our interpretation follows. We do not show any $N_{\text{fg}} < 4$ results as they are widely oscillating, namely, too much foreground contamination is still in the maps, and no measurement is possible. With the $N_{\text{fg}} = 4$ cleaned cube, the $P_{\text{Hr,g}}(k)$ eventually reaches a reasonable 'shape'; indeed, besides the amplitude, the consistency with CL23 indicates, first of all, a theoretically expected $P(k)$ behaviour. As contaminant suppression becomes more aggressive (increasing N_{fg}) the cosmological signal is increasingly suppressed as well, and the $P_{\text{Hr,g}}(k)$ gradually loses power. Eventually, for $N_{\text{fg}} = 6-7$, we are fully compatible with a null spectrum: the correlation is lost.

Next, we computed the cross- $P(k)$ of the mPCA-cleaned cubes. The results are shown in Fig. 13 where, with the same colour codes of Fig. 12, we show different cleaning with dark to light colour lines and areas (mean values and 1σ errors) and the red dashed line is the CL23 fitted model. For the mPCA case, we can choose how many modes to remove on the large (N_L) and – independently – on the small scale (N_S); the spectra in the figure have been derived varying N_L and keeping N_S fixed to 4 (top panel) and 8 (bottom). We find a similar – and actually clearer – trend as what we described above for Fig. 12: the cleaned cube that gives the highest $P_{\text{Hr,g}}(k)$ amplitude corresponds to a 'sweet spot' when varying the N_{fg} parameters. For brevity, we decided not to show results for varying N_S ; we find that the large scale N_L is the primary driver of the overall amplitude of the cross- $P(k)$.

7.2. Least-squares fitting

Next, we conducted a least-squares fit to the measured $P_{\text{Hr,g}}(k)$ with the model described in Sect. 6, and measure their amplitude $\Omega_{\text{Hr}}b_{\text{Hr}r}$. We assumed a Gaussian covariance, effectively ignoring correlation among adjacent k -bins⁸. We fit for one parameter, the amplitude of the cross- $P(k)$, and worked with 11 k -bins (data points), yielding 10 degrees of freedom. The results are in

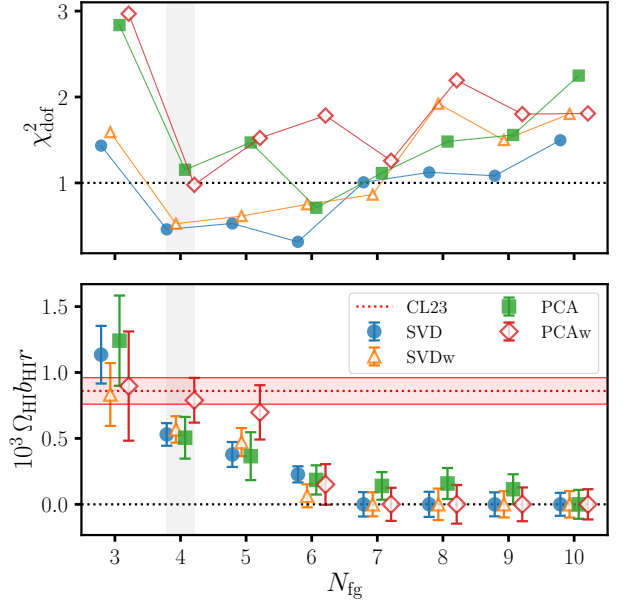


Fig. 14. Sensitivity of our results to the PCA-like contaminant subtractions as function of how many modes have been removed from the original cube. In the top panel, we show the reduced χ^2 for each foreground cleaned cross-power spectra relative to its best-fitting model; in the bottom panel, the corresponding best-fitting $\Omega_{\text{Hr}}b_{\text{Hr}r}$ values and their 1σ error bars. Filled symbols, circles and squares, are for an SVD cleaning respectively; their weighted counterparts corresponds to empty triangles and diamonds. In the bottom panel, the shaded red region highlights the CL23 best-fitting value and 1σ uncertainty that was obtained with a PCAw $N_{\text{fg}} = 30$ cleaning and a \mathcal{T} -correction (the 'uncorrected' CL23 amplitude would lay between $0.2-0.3 \times 10^{-3}$). We highlight in grey the $N_{\text{fg}} = 4$ results that correspond to our best and reference PCA result: the smallest N_{fg} to reach a plausible χ^2 .

Fig. 14 for the PCA cubes and Fig. 15 for mPCA as a function of the number N_{fg} of modes removed.

We start by looking at the PCA-framework results in Fig. 14. The top panel shows the reduced $\chi^2_{\text{d.o.f.}}$ of the best-fitting model; the bottom panel shows the fitted amplitudes with their 1σ uncertainties and, with the red dotted line and shaded area, the CL23 result and 1σ uncertainty. Different symbols and lines correspond to differently cleaned cubes: SVD (blue circles), PCA (green squares) and their weighted counterpart (empty triangles and diamonds for SVDw and PCAw respectively). Among those N_{fg} solutions in good agreement with the model ($\chi^2_{\text{d.o.f.}} \approx 1$), $N_{\text{fg}} = 4$ is the lowest N_{fg} and corresponds to the highest significance (highest amplitude) result, independently of the cleaning pipelines shown. The story told by Fig. 14 quantitatively matches our interpretation of the power spectra shown in Fig. 12: the power spectra are unrealistic for $N_{\text{fg}} < 4$, for $N_{\text{fg}} \geq 4$ the theoretical model can eventually describe them and we progressively see suppression of the power (signal) for $N_{\text{fg}} > 4$, finally reaching a null detection for $N_{\text{fg}} > 7-8$. Our best result – the best compromise between suppression of contaminants and suppression of the cosmological signal – is given by a PCAw analysis with $N_{\text{fg}} = 4$. We notice that PCA and PCAw give the lowest and highest amplitude among the four cleaned cubes. Nevertheless, all four of them are in agreement: they are less than 1σ away (2σ for PCA) the \mathcal{T} -corrected CL23 amplitude, and higher than the uncorrected CL23 measured amplitude ($\approx 0.25 \times 10^{-3}$), suggesting less signal loss in this analysis compared to CL23. Another point made by this plot is about weighting that improves the PCA and SVD solutions. Indeed, PCAw and SVDw

⁸ In ongoing work with higher signal-to-noise MeerKAT HI intensity mapping data, we are going beyond the Gaussian covariance approximation, e.g. MeerKAT Collaboration (2025).

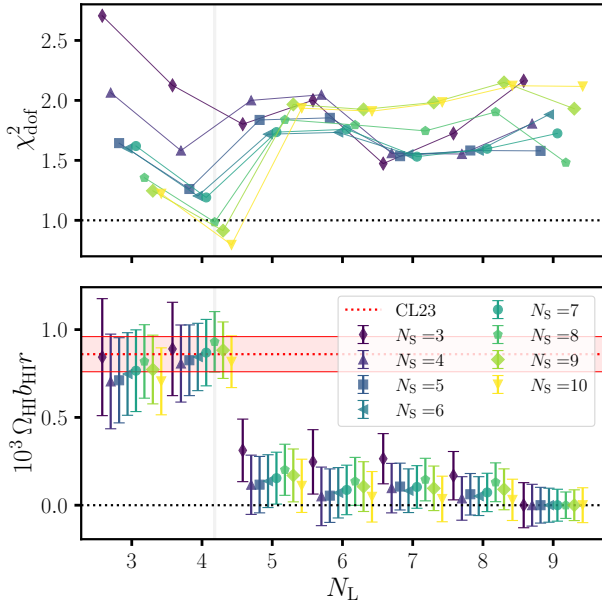


Fig. 15. Sensitivity of our results to the mPCA contaminant subtraction as function of how many modes have been removed from the original cube. The top panel shows the reduced χ^2 for each foreground cleaned cross-power spectra relative to its best-fitting model. The bottom panel the corresponding best-fitting $\Omega_{\text{HI}} b_{\text{HI}} r$ values and their 1σ error bars. All quantities are plotted as a function of the N_{fig} pair of the mPCA analysis, with N_L on the x -axis and N_S colour- and symbol-coded spanning from $N_S = 3$ in violet to $N_S = 10$ in yellow. For comparison, the shaded red region highlights the CL23 best-fitting value and 1σ uncertainty, that was obtained with a \mathcal{T} -correction (the ‘uncorrected’ CL23 amplitude would lay between $0.2\text{--}0.3 \times 10^{-3}$). We highlight in grey the $N_L, N_S = 4, 8$ results that correspond to our best and reference mPCA result. Results from mSVD and the weighted mSVDw and mPCA w show negligible differences.

(i) give more reasonable fits than PCA and SVD ($\chi^2_{\text{d.o.f.}}$ closer to 1, top panel), (ii) exhibit more clearly the ‘sweet’ spot at $N_{\text{fig}} = 4$ – it is unambiguously the minimum of their $\chi^2_{\text{d.o.f.}}$ curves – and, for those good fits, (iii) correspond to the highest amplitude of the cross-spectra. We recall here that although both signal loss and residual contaminants are expected in the cleaned cubes, we are not concerned about the latter in the cross-correlation, as those remaining contaminants should not correlate with the galaxy positions; hence, we can consider these measurements as lower limits and, thus, the highest amplitude is closest to the true one.

We now focus on the least-square fitting of the mPCA spectra; results are summarised in Fig. 15. The top panel shows the reduced $\chi^2_{\text{d.o.f.}}$ of the fits and the bottom the derived amplitude values and their standard deviations. On the x -scale, we show the number of coarse-scale removed modes N_L ; instead, different values of N_S correspond to different colours and symbols. We can confirm it is the N_L parameter that mainly drives the overall amplitude; however, we need to set also N_S accordingly to yield the best fits, in terms of the χ^2 value. For example, all $N_L = 4$ amplitude values lie in the expected range (CL23 result); yet, we need $N_S > 5$ to get a satisfactory χ^2 , with, for instance, $N_S = 3$ always leading to an arguably high χ^2 regardless of the N_L we set. We can spot another interesting trend: the $N_S = 8$ cases maximise the amplitude for any given N_L . We can speculate that here N_L plays the most important role because we are fitting for an amplitude rather than a shape, and indeed N_S is important for reaching a good χ^2 . Hence, the small-scale cleaning will become

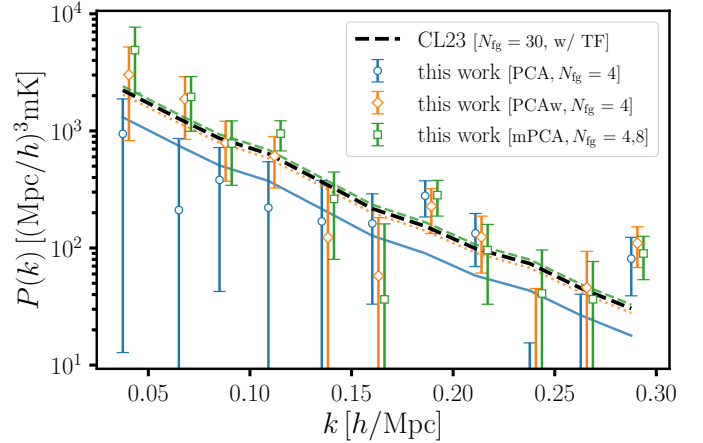


Fig. 16. Cross-power spectra between WiggleZ galaxies and MeerKAT HI intensity maps at redshift $z \approx 0.4$. Maps have been cleaned by a PCA and PCAw analysis with $N_{\text{fig}} = 4$ (blue circles and orange diamonds respectively) and by mPCA with $N_L, N_S = 4, 8$ (green squares). We show error bars with 1σ uncertainties and the resulting best-fit curves (solid blue for PCA, dotted orange for PCAw and dashed green for mPCA). The black dashed line is the fitted model from CL23 (same data, different cleaning strategy) that instead was enhanced by a transfer function correction.

crucial when data quality will allow us to constrain cosmology, namely, the shape of the power spectrum.

As we did for the PCA-like methods, among the cubes with an acceptable χ^2 , as the mPCA reference, we selected the cube with the highest amplitude (and, hence a higher detection significance), which is given by the $N_L, N_S = 4, 8$ combination.

We stress again that, for all the cleaning strategies reported in Figs. 14 and 15, the best-performing residuals (good fit and highest amplitude) correspond to values of $\Omega_{\text{HI}} b_{\text{HI}} r$ that are in full agreement with the \mathcal{T} -corrected CL23 result. Yet, their statistical uncertainties are higher than the 12% for CL23: 30% for PCA, 21% for PCAw and 18% for mPCA. However, CL23 (i) performed the extra ‘reconvolution’ smoothing, suppressing power, and (ii) removed more components (30 versus 4) suppressing thermal noise⁹; both decrease the cubes auto-spectra, influencing the Gaussian errors’ calculations. Regardless, CL23 experienced a degeneracy with the number of removed modes – it was less clear which N_{fig} to set, leading to an additional variance of 14%. Interestingly, when we account for this error in the overall $\Omega_{\text{HI}} b_{\text{HI}} r$ uncertainty, we find it comparable with what is estimated in the present analyses.

We summarise the results from all contaminant-separation pipelines in Table 1, and we further showcase the PCA, PCAw and mPCA solutions in Fig. 16, with the results from CL23 for comparison. Table 1 displays the cleaning level (N_{fig} or N_L, N_S), whether or not a transfer function correction is in place (\mathcal{T}), the measured amplitude of the cross-spectra ($\Omega_{\text{HI}} b_{\text{HI}} r$), the goodness of the fit ($\chi^2_{\text{d.o.f.}}$), and its statistical detection significance

$\sqrt{\Delta\chi^2} \equiv \sqrt{\chi^2_{\text{null}} - \chi^2}$, namely, the difference between the χ^2 evaluated using our cross-correlation model, and one using a null model with zero power. The uncertainty that we quote for the CL23 amplitude is the sum in quadrature of the two reported errors ($\Omega_{\text{HI}} b_{\text{HI}} r = [0.86 \pm 0.10(\text{stat}) \pm 0.12(\text{sys})] \times 10^3$); its significance has also been updated accordingly.

⁹ Uncorrelated noise suppression resembles signal loss; both are filtered similarly due to their similar frequency structure.

Table 1. Results of the least-squares fits of the cross-power spectra of the MeerKAT HI IM maps and the WiggleZ galaxy at redshift $z \approx 0.4$.

Method	N_{fg}	\mathcal{T}	$\Omega_{\text{HI}} b_{\text{HI}} r \times 10^3$	$\chi_{\text{d.o.f.}}^2$	significance
SVD	4	no	0.63 ± 0.13	0.78	3.9σ
SVDw	4	no	0.75 ± 0.16	0.91	4.4σ
PCA	4	no	0.50 ± 0.16	1.15	3.3σ
PCAw	4	no	0.79 ± 0.17	0.98	4.7σ
mSVD	4, 8	no	0.91 ± 0.17	0.99	5.6σ
mSVDw	4, 8	no	0.86 ± 0.17	1.01	5.3σ
mPCA	4, 8	no	0.93 ± 0.17	0.99	5.8σ
mPCAw	4, 8	no	0.87 ± 0.17	0.97	5.4σ
CL23	30	yes	0.86 ± 0.16	0.84	4.8σ

Notes. Each row corresponds to a different cleaning strategy. Being limited by signal loss, the highest amplitudes are the closest-to-truth, although all amplitude values agree within $\leq 2\sigma$. In bold, we highlight: PCA and PCAw to bracket the PCA-like results, mPCA as representative of the multiscale solutions and our benchmark from CL23 (last row), which corresponds to a weighted PCAw with different preprocessing choices and a transfer function correction applied.

7.3. Null tests and jackknife errors

To validate the detections discussed above, we perform null tests by randomising galaxies along the frequency direction, see Fig. 17. Each panel refers to a different reference cube (from left to right: PCA $N_{\text{fg}} = 4$, PCAw $N_{\text{fg}} = 4$ and mPCA $N_L, N_S = 4, 8$), whose cross- $P(k)$ is displayed with orange dots. The 100 cross- $P(k)$ with the 100 randomised galaxy fields are in light blue, their overall mean and standard deviation correspond to the empty diamonds: fully compatible with zero.

We believe that the uncertainties of our measured $P_{\text{HI,g}}(k)$ are adequate as the reasonable $\chi_{\text{d.o.f.}}^2$ of the fitted reference $P_{\text{HI,g}}(k)$ indicate. However, we further validate them by comparing them with jackknife errors (Norberg et al. 2009). We divide the cleaned cubes into 320 N_{jk} equal sub-cubes (using a higher N_{jk} does not produce worthwhile different results) and compute the 320 cross-spectra $P_{\text{jk}}^i(k)$. The final jackknife spectrum is given by the mean $\mu_{\text{jk}}(k)$ of the N_{jk} spectra, and its uncertainties are computed as follows (Hartlap et al. 2007; Favole et al. 2021):

$$\sigma_{\text{jk}}^2(k) = \frac{N_{\text{jk}} - 1}{N_{\text{jk}}} \frac{N_{\text{jk}} - 1}{N_{\text{jk}} - N_k - 2} \sum_{i=1}^{N_{\text{jk}}} \left[P_{\text{jk}}^i(k) - \mu_{\text{jk}}(k) \right]^2, \quad (17)$$

with $N_k = 11$ the number of k -bins. We plot the $\{\mu_{\text{jk}}(k), \sigma_{\text{jk}}(k)\}$ data points in the panels of Fig. 17 with filled green triangles to compare them with our reference spectra in orange. As expected, jackknife errors tend to be higher because of the non-independence of the jackknife sub-cubes. However, overall we find a satisfying agreement with the Gaussian errors of our reference analysis. Indeed, the best-fit of $\{\mu_{\text{jk}}(k), \sigma_{\text{jk}}(k)\}$ is almost indistinguishable from the reference one (dashed green versus solid orange lines), with the – slight – exception of PCAw (middle panel). Being more quantitative, compared to the reference analyses, the higher $\sigma_{\text{jk}}^2(k)$ reduce the $\chi_{\text{d.o.f.}}^2$ by $\sim 40\%$ for all three reference cubes considered, and lead to practically unvaried amplitudes and associated error values, with the exception of PCAw: $\Omega_{\text{HI}} b_{\text{HI}} r \times 10^3 = 0.84 \pm 0.19$ compared to the reference 0.79 ± 0.17 .

7.4. Scale-independence of the fitted spectra

This paper focuses on reaching a compromise between obtaining a good contaminant subtraction while minimising cosmological signal loss. Hence, a test of the robustness of our cross- $P(k)$ measurements is to check their independence with respect to scale; namely, we want to check whether the choice of relying on either large-scale or small-scale information results in significant changes. Indeed, if a certain k -scale is more prone to signal loss (typically the large scale), the least-square fit might lead to different solutions when used on a subset of the $P_{\text{HI,g}}(k)$ sampled k -bins.

We tested the scale-independence of the reference cross-spectra obtained with PCA, PCAw, and mPCA, with the results shown in the three columns of Fig. 18 (from left to right). For each spectrum, we perform again the least-square fit discarding 1, 2 or 3 k -bins either at the low or high end of the available k -range; namely, we produce six new fits for each measured $P_{\text{HI,g}}(k)$, neglecting the large or small scale. For each of these new fits, we can associate an effective scale

$$k_{\text{eff}} = \frac{\sum_i^{N_k} k_i (P(k_i)/\sigma(k_i))}{\sum_i^{N_k} (P(k_i)/\sigma(k_i))}. \quad (18)$$

When discarding low (high) k -bins k_i , k_{eff} grows (lowers). In all panels of Fig. 18, we highlight with a vertical line the k_{eff} of the reference spectra, using all 11 k_i bins: for each column of Fig. 18, we have three data points on the left, three on the right of that line, for the large-scale and small-scale-only fits respectively. We notice that the PCA results (left panels) are in the right area of the plot: the signal-to-noise $P(k_i)/\sigma(k_i)$ is generally worse for low k_i compared to mPCA (right panels); the PCAw case lies between the two (middle panels).

In the top panels of Fig. 18, we show the $\chi_{\text{d.o.f.}}^2$ of the fits as a function of k_{eff} . In all three scenarios, the low k_i bins are associated with higher relative uncertainties, hence the large-scale fits tend to display lower $\chi_{\text{d.o.f.}}^2$ than the small-scale fits. Anyway, these new fits show good $\chi_{\text{d.o.f.}}^2$: the cross-correlation theoretical model describes data reasonably well regardless of the k -scale we analyse. The latter is especially true for the mPCA case (last column).

In the bottom panels of Fig. 18, we compare the fitted amplitudes of $\Omega_{\text{HI}} b_{\text{HI}} r$, that for brevity we call A^i , against the reference $A^{\text{ref}} \pm \sigma_A^{\text{ref}}$, namely, obtained with the all k -bins fit reported in the text inside each panel—through the quantity:

$$\Delta_{\text{amp}}^i = \frac{A^i - A^{\text{ref}}}{\sigma_A^{\text{ref}}}. \quad (19)$$

With Δ_{amp} , we quantify the agreement of the new fits of the robustness test in units of the uncertainty σ_A^{ref} . All the 6×3 fitted amplitudes are within 1σ from their reference values and, although not shown, with highly similar detection significances—especially true for PCAw and mPCA: the agreement is remarkable, supporting our detection measurements. By contrast, in CL23, the $\Omega_{\text{HI}} b_{\text{HI}} r$ amplitude was changing by a factor of ~ 2 (with corresponding $\Delta_{\text{amp}} > 1\sigma$) in a similar k range (Figure 7 therein), leading to the possibility of the presence of spurious signal in the large scale or a k -dependence unexpectedly introduced by the \mathcal{T} -correction. In the analysis of this work, when we remove the large scale, results do not change, and even more impressively, when we mainly rely on the large scale, by excluding the small-scale, we still get the same results.

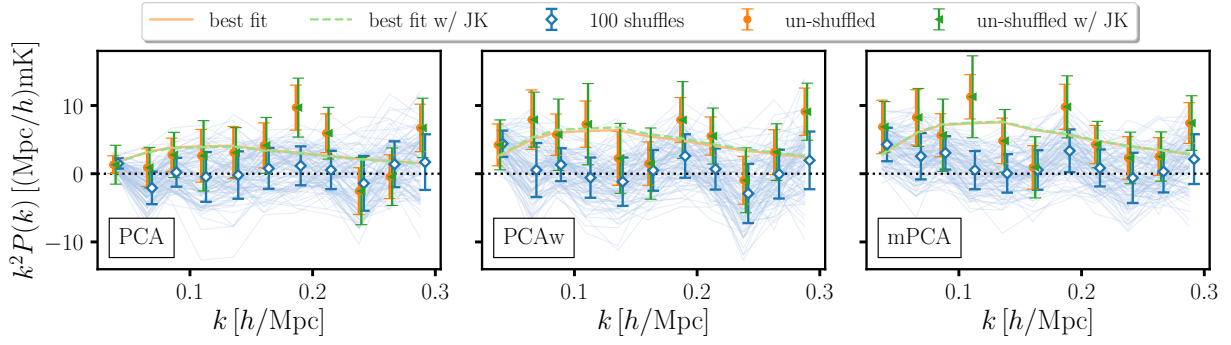


Fig. 17. Null tests for the cross-power spectra derived with the PCA $N_{\text{fig}} = 4$ (left panel), its weighted version PCAw $N_{\text{fig}} = 4$ (middle) and multiscale version mPCA $N_L, N_S = 4, 8$ (right). The WiggleZ galaxy maps have been shuffled along frequency (redshift). The thin blue lines show the spectra computed for 100 different shuffles. Empty blue diamonds correspond to the average and standard deviation across the shuffled samples, that we compare against (i) the reference spectra (filled orange circles) and (ii) the reference spectra with jackknife errors (filled green triangles). We use the solid orange and dashed green lines to show the best fits of the reference spectra, using Gaussian and jackknife errors respectively.

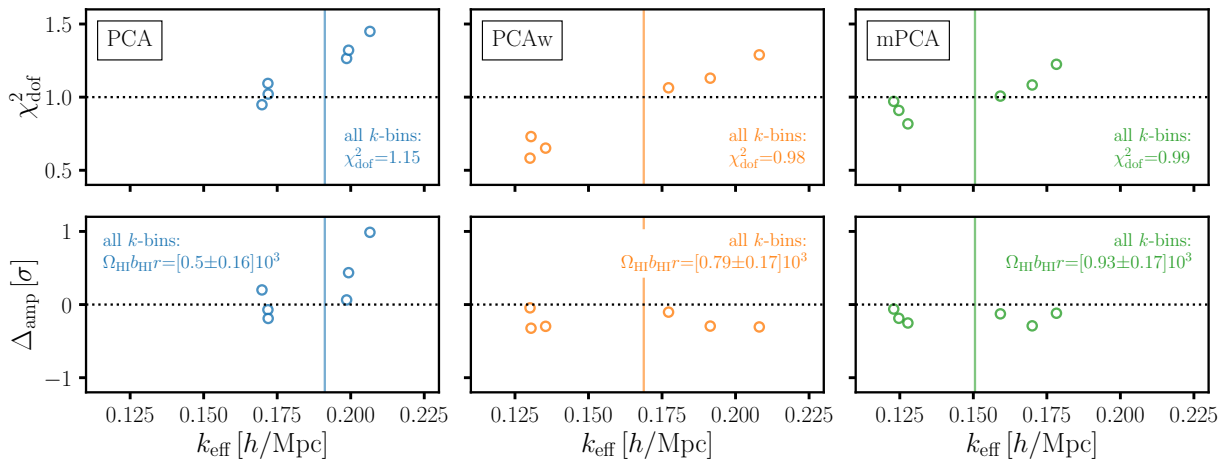


Fig. 18. Robustness of the cross-spectra fits. From left to right, we test the scale-independence of the fits for the measured cross- $P(k)$ between the WiggleZ galaxies and the PCA $N_{\text{fig}} = 4$ cleaned cube (left column), PCAw $N_{\text{fig}} = 4$ (middle) and mPCA $N_L, N_S = 4, 8$ (right). *Top panels:* Reduced χ^2 for the least-squares analysis as function of the effective scale of the data cube k_{eff} when excluding from 1 to 3 smallest and highest k -bins from the fit, namely, by removing either large or small-scale information. *Bottom panels:* Agreement of the amplitudes of the less-information fits with respect to the reference all- k -bins fitted amplitude, through the quantity Δ_{amp} , that quantifies how many σ away the $\Omega_{\text{HI}} b_{\text{HI}} r$ values are from the reference, indicated in the text for each panel. In all panels, the vertical line shows the corresponding k_{eff} for the reference all- k -bin analysis for each cleaning method.

7.5. Assessing leakages in the component separation

Our starting point for subtracting contaminants from the MeerKAT intensity maps is the linear decomposition in Eq. (4). Contaminant separation aims to minimise leakages between the sum terms of Eq. (4).

For each of the cleaned, reference cubes of this work, we statistically evaluate the leakages between the summed terms of Eq. (4) by cross-correlating $\mathbf{X}_{\text{clean}}$ with the estimated foregrounds $\mathbf{X}_{\text{FGs}} = \mathbf{X} - \mathbf{X}_{\text{clean}}$. Such a signal could be produced if the contaminants are insufficiently characterised and leak into the residuals or if the removed foregrounds contain some cosmological signal (i.e. signal loss). We compare these cross-spectra with their respective autos through the cross-correlation coefficient, r_{LK} :

$$r_{\text{LK}} = \frac{\mathcal{P}(\mathbf{X}_{\text{clean}}, \mathbf{X}_{\text{FGs}})}{\sqrt{\mathcal{P}(\mathbf{X}_{\text{clean}}, \mathbf{X}_{\text{clean}})\mathcal{P}(\mathbf{X}_{\text{FGs}}, \mathbf{X}_{\text{FGs}})}}, \quad (20)$$

where \mathcal{P} denotes the power spectrum computation described in Sect. 6 (e.g. following this notation, the cross-spectra earlier measured in this paper would be expressed as $\hat{P}_{\text{HI},g} \equiv$

$\mathcal{P}(\mathbf{X}_{\text{clean}}, \mathbf{X}_{\text{galaxies}})$). We show the behaviour of r_{LK} for the different cleaned cubes in Fig. 19. The correlation between the estimated foregrounds and the estimated cosmological signal oscillates around zero, with a maximum, absolute amplitude of $|r_{\text{LK}}| < 0.15$ at $k > 0.06 h/\text{Mpc}$ for all three (PCA, PCAw, and mPCA) solutions, highlighting a reasonable degree of no-leakage¹⁰.

7.6. Assessing signal loss: Mock-injection tests

The contaminant separation methods filter out some cosmological signal from the final cleaned maps. This effect is particularly significant on large scales, leading to a suppression of the signal power spectrum. Hence, through mock-injection tests, a corrective transfer function \mathcal{T} can be constructed. The \mathcal{T} has been widely used in the analyses of HI intensity mapping observations

¹⁰ The exception is the first k -bin sampled that corresponds to $0.3 < r_{\text{LK}} < 0.5$. However, we recall that when we discard this largest k -bin measurement from the final fit, our analysis yields amplitude measurements statistically unchanged (Fig. 18).




Fig. 19. Cross-correlation coefficient between reconstructed foregrounds and cosmological signal, as defined in Eq. (20). We show results for the PCA $N_{\text{fg}} = 4$ cube (solid blue line), PCAw $N_{\text{fg}} = 4$ (dotted orange) and mPCA $N_L, N_S = 4, 8$ (dashed green).




Fig. 20. *Top panel:* Transfer function $\mathcal{T}(k)$ of the three reference analyses: PCA $N_{\text{fg}} = 4$ (solid blue), PCAw $N_{\text{fg}} = 4$ (dotted orange) and mPCA $N_L, N_S = 4, 8$ (dashed green) and their 1σ scatter (over 200 mocks) with corresponding shaded areas. *Bottom panel:* Signal-to-noise ratio of $\mathcal{T}(k)$, i.e. the ratios over the 1σ scatter.

(see e.g. Masui et al. 2013; Anderson et al. 2018; Wolz et al. 2022; Cunnington et al. 2023a, and Switzer et al. (2015) for a more in-depth discussion). In practice, it enhances the power spectrum to counteract the estimated suppression.

Although we have demonstrated that the cross-power spectrum measurements in this work, given their uncertainties, do not significantly suffer from signal loss because we do not see any evidence of its k -dependent effect, we further corroborate that statement by computing the \mathcal{T} function and correcting the measurements accordingly.

We followed the recipe derived by Cunnington et al. (2023b), who also validated a method that uses the transfer function variance for error estimation for the power spectra measurements. We began by generating N_m mock HI signal cubes \mathbf{X}_m , in boxes of same size and gridding as our observed cube; we followed the lognormal approximation (Coles & Jones 1991) and sampled from the theoretical power spectrum model in Eq. (14), with cosmology and beam parameters as in our fitting proce-

cedure, and HI bias and mean brightness temperature values of CL23, assuming $r = 1$. Each mock is then injected into the actual data. We run our component separation pipelines (PCA, PCAw, or mPCA) on the combination $\mathbf{X} + \mathbf{X}_m$. Thus, for mPCA, this process also entails going through the wavelet-filtering and back. In short, for each $\mathbf{X} + \mathbf{X}_m$ and each method, we determine the mixing matrix $\hat{\mathbf{A}}_m$ that yields the cleaned cube,

$$\mathbf{X}_{\text{clean}}^m = (\mathbf{X} + \mathbf{X}_m) - \hat{\mathbf{A}}_m (\hat{\mathbf{A}}_m^T \hat{\mathbf{A}}_m)^{-1} \hat{\mathbf{A}}_m^T (\mathbf{X} + \mathbf{X}_m) - [\mathbf{X}_{\text{clean}}]. \quad (21)$$

The last term in the square brackets is the original cleaned observed data, which we subtract when we want to remove contributions in the cube uncorrelated to the mock signal, thus reducing the variance of the \mathcal{T} function. In practice, we subtract it when computing the $\mathcal{T}(k)$ but not when computing its covariance. The transfer function is defined as:

$$\mathcal{T}(k) = \left\langle \frac{\mathcal{P}(\mathbf{X}_{\text{clean}}^m, \mathbf{X}_m)}{\mathcal{P}(\mathbf{X}_m, \mathbf{X}_m)} \right\rangle_{N_m}, \quad (22)$$

where \mathcal{P} denotes the power spectrum computation described in Sect. 6; the external brackets in Eq. (22) correspond to taking the median over N_m iterations, one for each mock.

Applying the transfer function correction on the measured $P(k) = \mathcal{P}(\mathbf{X}_{\text{clean}}, \mathbf{X}_{\text{galaxies}})$ boils down to the following:

$$P_{\text{TF}}(k) = \frac{P(k)}{\mathcal{T}(k)}, \quad (23)$$

with uncertainties of

$$\sigma[P_{\text{TF}}(k)] = \sigma \left\{ \frac{P(k)}{\mathcal{T}_i(k)} \right\}, \quad (24)$$

where \mathcal{T}_i is the transfer function from the i^{th} mock iteration and with $\sigma\{\}$ we take the scatter over all iterations.

Results are in Fig. 20, for PCA in solid blue, wPCA in dotted orange and mPCA in dashed green, with the shaded areas corresponding to their 1σ scatter. The behaviour of the three \mathcal{T} functions is as expected: very pronounced – high eventual correction – at the smallest k and quickly increasing to approach 1 at the greatest k . It is reassuring to see that the pipelines' suppression levels are in the same order as their cross-spectra amplitudes, namely, the lowest measured amplitude value is associated with the highest signal suppression in simulations. Different signal loss levels are likely the reason for the differently cleaned cubes to lead to different cross-spectrum amplitudes.

What is striking in Fig. 20 is the amount of variance related to the transfer functions. The plot shows results using $N_m = 200$ mocks; they converge beyond this number. To highlight the \mathcal{T} -driven uncertainties, we show the signal-to-noise ratio, $\mathcal{T}/\sigma_{\mathcal{T}}$, in the bottom panel of Fig. 20, which is below unity for the first half of the k -bins we sample and hardly reaches 2–3 at the smallest scale. Understanding the reason why the \mathcal{T} calculations are so undetermined in our analyses is beyond the scope of the present work. Recently, Chen (2025) used a quadratic estimator approach and simulations to demonstrate that, in the presence of mode-mixing, the \mathcal{T} correction is (i) potentially biased and (ii) leads to an incorrect covariance estimation. Here, we have the chance to test these findings with observations.

We proceed to correct the measured spectra with \mathcal{T} as in Eq. (23) and assign errors following Eq. (24); we re-run the least-square fitting pipeline and determine the best-fitting amplitude values for the different spectra $P_{\text{TF}}(k)$. We show results in Fig. 21. As in previous similar figures, we show in the top the goodness of the fits, and in the bottom the fitted amplitudes, both

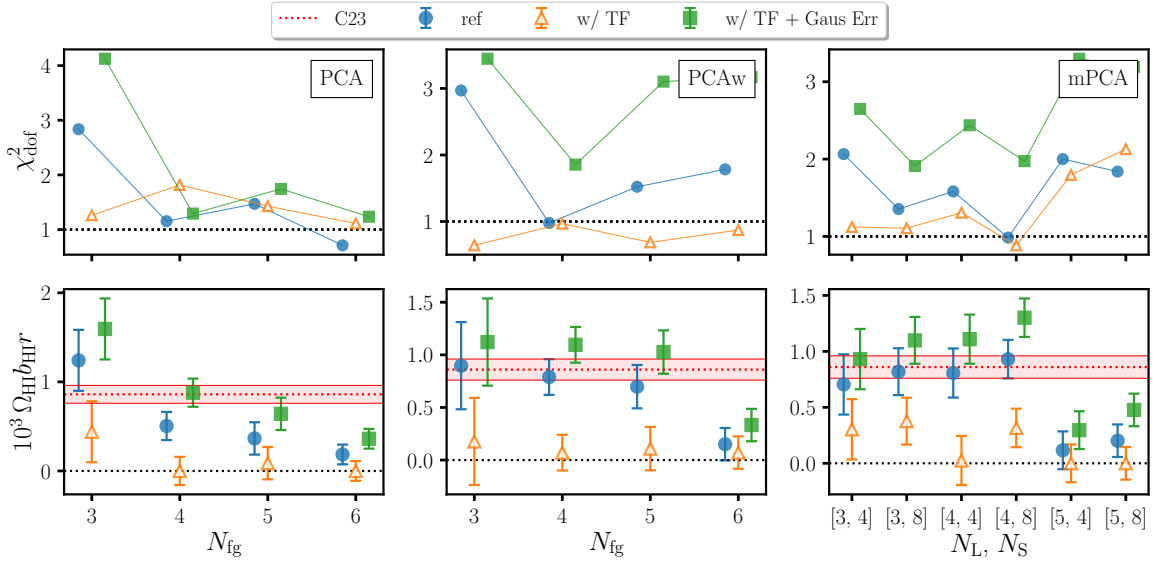


Fig. 21. Top panels: Reduced $\chi^2_{\text{d.o.f.}}$ of the cross-power best fits. Bottom panels: Corresponding amplitude values compared to the CL23 result in red, both as function of the cleaning level N_{fg} . Different columns correspond to different methods: PCA, PCAw and mPCA from left to right. The filled blue circles correspond to the ‘standard’ reference analysis so far discussed; the empty orange triangle to the analysis performed with the transfer function correction; the filled green squares correspond to the analysis performed with the transfer function correction and imposing Gaussian uncertainties for the final, corrected measurements.

as a function of the number N_{fg} of modes removed. The outcome of the $P_{\text{TF}}(k)$ fits are marked with empty orange triangles, to be compared to the standard analysis (no \mathcal{T} -correction) with filled blue circles. Results are equivalent for all methods (PCA, PCAw and mPCA from left to right): uncertainties are so large that all $P_{\text{TF}}(k)$ are in agreement with a zero correlation.

Hypothesising that something is misguided with the uncertainties only (one of the findings of Chen 2025) and that the \mathcal{T} values themselves are reliable estimates of the clustering suppression, we perform the correction as in Eq. (23), but without using the \mathcal{T} -variance driven errors; instead, we arbitrarily apply Gaussian statistics to the uncertainties and use Eq. (11) for the $P_{\text{TF}}(k)$ measurements, substituting $\hat{P}_{\text{HI,g}}$ and \hat{P}_{HI} with the enhanced, \mathcal{T} -corrected spectra. Results of this latter exercise are shown as green squares in Fig. 21. The only case for which fits are as reasonable as for the uncorrected spectra is the PCA analysis, in the left column. Interestingly, the new $N_{\text{fg}} = 4$ amplitude has increased and is now in better agreement with CL23 and the analyses of other methods. However, the transfer function correction should work for all PCA analyses with different N_{fg} that result in differently suppressed signal clustering, and all the corrected $P_{\text{TF}}(k)$ should be unbiased, that is, pointing to the same reconstructed amplitude. However, we do not see this for the PCA $N_{\text{fg}} = 4, 5$, and 6 solutions, which continue to decrease as function of N_{fg} , like the uncorrected, standard values (blue circles). Finally, for PCAw and mPCA, these last fits’ goodness has worsened compared to the references; in terms of spectra amplitudes, results are less than 1σ away from the reference values, in agreement with the claim that, given our uncertainties, these cubes do not experience important signal loss effects.

In summary, our computed transfer functions are biased and linked to unreliable covariances, supporting the theoretical work by Chen (2025). We likely still lack a method to compute \mathcal{T} correctly and propagate uncertainties in this data regime; namely, low signal-to-noise ratios and high corruption by systematics that lead to mode-mixing. However, we do not want to dedicate efforts to strengthening the \mathcal{T} computation since (i) this work is focussed on comparing the cleaning methods and the most

optimal pre-processing choices they require and (ii) we are currently forward-modelling the signal loss effect in the upcoming MeerKLASS observations’ analyses, finding it a viable method since data quality and signal-to-noise ratios are steadily improving, leaving constraining power for additional nuisance parameters.

8. Peering towards an independent hydrogen intensity mapping detection with the mPCA cleaning

8.1. Assessing the cubes’ variance in retrospect

As discussed earlier, in a PCA framework, the quicker (with a lower N_{fg}) we can reduce the variance of the data cube, the more efficient the decomposition. If removing an additional component does not visibly change the variance, the assumption of PCA itself is no longer valid, as the principal components should ‘explain’ variance.

In retrospect, after having measured the cross-spectrum with the external tracer, we can compare the best-cleaned cubes’ variances. We do so in Fig. 22, yielding counter-intuitive results. The cube that retained most of the cosmological signal – higher cross-correlation amplitude – is not that with the overall highest variance, namely, the red, transparent horizontal stripe of mPCA, compared to PCA (blue) and PCAw (orange). Hence, the higher variance of PCA and PCAw is likely linked to higher contaminant residuals in the maps than mPCA. The latter hypothesis is supported by the robustness tests’ results shown in Fig. 18. Foreground residuals will become problematic as we independently measure the HI intensity mapping signal: in such cases, we would need more aggressive cleaning, further worsening the signal loss problem.

In Fig. 22, we also plot the $N_S = 4$ mPCA solutions as a function of N_L , to highlight how, although not fully appreciable in the cross-correlation measurement, N_S plays a crucial role in the mPCA decomposition.

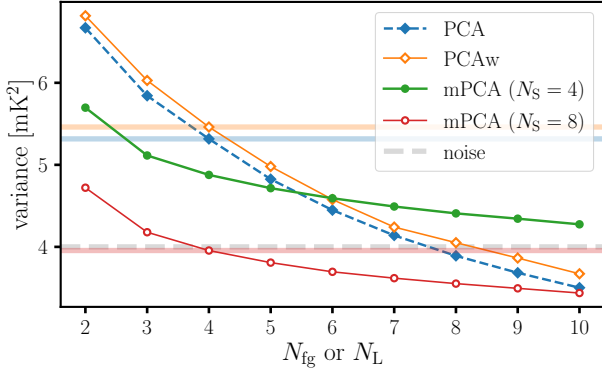


Fig. 22. Variance in the cleaned cubes as a function of the number of components removed: N_{fg} or N_L for PCA (dashed blue line), PCAw (solid orange), and mPCA (solid green setting $N_S = 4$ and solid red for $N_S = 8$). Horizontal stripes with corresponding colours highlight the variance level of the reference cube for each method ($N_{fg} = 4$ for PCA and PCAw and $N_L, N_S = 4, 8$ for mPCA), namely the solutions for each method that offer a maximised goodness-of-fit and amplitude when cross-correlated with the galaxy catalogue. For reference, the variance of the original, unclean cube corresponds to $\approx 9 \times 10^3 \text{ mK}^2$ and is characterised by the noise level highlighted by the dashed gray stripe.

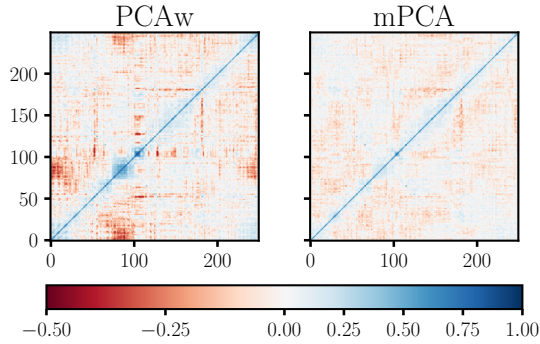


Fig. 23. Frequency-frequency correlation matrices of the cleaned cubes, PCAw with $N_{fg} = 4$ on the left and mPCA with $N_L, N_S = 4, 8$ on the right. Zero correlation corresponds to the white colour. Both matrices show the expected positive diagonal structure of the HI IM signal. The mPCA solution results in fewer off-diagonal contributions.

8.2. Correlation matrices: Off-diagonal terms

We plot the normalised $\mathbf{X}_{\text{clean}} \mathbf{X}_{\text{clean}}^T$ matrices in Fig. 23. On the left panel, we show the solution of the PCAw analysis with $N_{fg} = 4$; on the right panel, the mPCA solution with $N_L, N_S = 4, 8$. Both matrices refer to cleaned cubes that show positive and comparable correlations with the WiggleZ catalogue. Yet, they look different: the matrix on the right (mPCA) is substantially ‘cleaner’ than its PCAw counterpart: the diagonal is positive, as expected, and with a much stronger correlation than any off-diagonal term. The secondary tracer—the galaxies—allowed us not to worry about the spurious correlations. We expect to start worrying about them as we proceed toward an independent clustering detection with the HI temperature fluctuations maps.

9. Discussion

9.1. Confirmation of the detection in cross-correlation

The current study enhances the CL23 detection of the cross-correlation signal between the MeerKAT HI intensity maps and

the WiggleZ galaxies at about $z \approx 0.4$. We summarise the cross- $P(k)$ detections, significances and fitted amplitudes in Table 1.

Although the data are the same as in CL23, pre-processing steps (fewer cuts in frequency space, more in pixel space, no reconvolution) and cleaning method details, especially for the novel multiscale component separation analysis, have changed. Actually, we used the cross-correlation of CL23 as a benchmark to test our signal-processing choices, and the exercise has led to improvements in our analysis pipeline, as we discuss below. Moreover, in the present analysis, all correlation disappears after $N_{fg} \approx 8$ components removed, compared with the optimal $N_{fg} = 30$ in CL23, a further indication that we have made contaminant subtraction more efficient.

Correlating the MeerKAT maps with the galaxy catalogue allows us to be moderate with the suppression of contaminants in the data. The aim is to find the best compromise between cleaning the maps without removing too much cosmological information. The sense of ‘compromising’ is clear if we look back at Figs. 12 and 14 for the standard PCA-like methods and Figs. 13 and 15 for the multiscale case. The $N_{fg} = 4$ and $N_L, N_S = 4, 8$ results are the best compromise to, on the one hand, remove enough contamination to ‘reveal’ a cosmological signal that correlates with the galaxies and, on the other hand, limit the cosmological signal loss. Indeed, for $N_{fg} < 4$, no detection is possible (too much contamination still present); for $N_{fg} > 4$, we have detection (up to $N_{fg} \approx 6$) but with gradually worse fits and less significance: we are experiencing cosmological signal loss. Eventually, we remove too much with $N_{fg} > 6$ and see nothing.

A welcome consequence of the situation described above is that, in our case, the optimal choice of N_{fg} is well motivated by the analysis: the chosen N_{fg} is the one that maximises the cross- $P(k)$ amplitude and gives plausible fits with the theoretical model. Indeed, the need to (almost) arbitrarily set N_{fg} has been a drawback of these component separation methods¹¹. This approach should work also when analysing HI data only, looking for a direct detection through internal cross-correlations (e.g. among time-split subsets of data), as the MeerKLASS collaboration plans to do in the future.

We corroborated the robustness of the cross- $P(k)$ detection with the goodness of the fit, the comparison with a zero-detection, null tests, and a demonstrated independence from k -scale cuts. Yet, what about the amplitude of the measured cross- $P(k)$? How solid are the HI properties we derived through the $\Omega_{\text{HI}} b_{\text{HI}}$ parameters, and why, compared to CL23, we did not correct for signal loss?

9.2. Beyond the state-of-the-art

The cross-power spectra of this work have not been boosted with a transfer function, \mathcal{T} , correction, yet they fully agree with CL23, where a strong \mathcal{T} was in place. This agreement is remarkable and has significant consequences. Most importantly, the signal loss in the cleaned cubes of this work is small enough to be negligible within the uncertainties that still characterise our analysis, as explicitly shown in Sect. 7.6 and confirmed by the robustness test on the k -independency of our results (Fig. 18).

We also seek to discuss why the resulting cross spectra of this work do not need a correction for power loss. Firstly, as pointed

¹¹ There have been propositions in the literature to set N_{fg} automatically (e.g. Olivari et al. 2016). However, these schemes typically need a prior for the signal, which is hard to provide for first-of-their-kind measurements, as the large scale HI IM signal we aim for, but could eventually become advantageous with the maturity of the field.

out in Sect. 3, the amount of cosmological signal lost during the component separation process grows with the number N_{fg} of components removed. Hence, keeping N_{fg} as low as possible is the most straightforward and orderly way to limit signal loss. Removing one order of magnitude fewer modes in this work than CL23 translates into preserving more signal.

Moreover, we expect the eventual correction to be k -dependent as we typically lose more signal on the large scale—it is where signal and foregrounds overlap the most, making them difficult to disentangle. This strong k -dependency is supported by analyses of different observations (Anderson et al. 2018; Wolz et al. 2022, CL23) and simulations alike. Yet, looking at the least-squares fits in Fig. 18 reveals that there is no possibility of a \mathcal{T} -correction to significantly alter the small and large k contributions, given the current measurement uncertainties.

In conclusion, the analysis of this work has further corroborated the HI properties that we infer from the MeerKAT L -band intensity maps of the 2019 Pilot Survey (as in Table 1). It has also gone beyond that since it has demonstrated that we can measure the cross-correlation signal relying on our novel contaminant suppression pipeline that adequately preserves the cosmological signal, is robust to scale-cuts and points to an unambiguous choice of the level of cleaning (N_{fg} parameter). All the latter features are novel in the HI intensity mapping observational context and represent a significant leap beyond the state-of-the-art.

9.3. Investigating what choices lead to the improvement

Compared to CL23, (i) we are more cautious in pixel space, reducing the footprint OF to the CF by analysing the components PCA determines when working in the OF and discarding roughly 60% of the pixels. However, here, (ii) we use the whole 250-channel range in the entire band (CL23 removed about one-third of the total frequencies at an additional RFI-flagging stage before the cleaning). To summarise, we ended up working with a comparable amount of data; the main difference is the choice of being more severe with the flagging in pixel rather than frequency space, making our cube ‘coherent’ and facilitating the PCA decomposition. Another critical difference is that (iii) we do not reconvolve the maps with a frequency-dependent Gaussian kernel, effectively not enhancing mode-mixing in the observed cube. Finally, (iv) we introduced a novel multiscale cleaning framework that has exceeded the performance of the other pipelines. Below, we summarise the fundamentals of this work’s contaminant subtraction analysis.

Handling outliers in pixel space. The weighting scheme based on the scan hits map is insufficient for taking care of the pixel outliers (i.e. pixels largely contributing to increasing the temperature variance of each map) that make the PCA decomposition deviate¹². At present, we find it is more efficient to flag those pixels. Indeed, we demonstrated the importance of performing the signal separation task in the smaller CF: the PCA solution substantially changes; a flagging in the pixel space done after the separation process does not have the same effect.

¹² In the presence of strong outliers in pixel space, PCA ‘wastes’ components to describe those features, and we have demonstrated that down-weighting these pixels does not seem enough to stop this from happening (similar conclusions are also drawn in, e.g. Wolz et al. 2017). Novel and more efficient weighting schemes might be possible too.

Coherence in the frequency direction. On the other hand, being too aggressive in the channel flagging is detrimental: having discontinuous information in frequency harms the PCA decomposition, which can describe abnormal variations in the frequency direction variations but struggles to detect the artificial ‘jumps’ in the spectral behaviour of the foregrounds. Moreover, dimensionality reduction techniques as PCA work best in the regime in which components are outnumbered by observations. Hence, we should try to maximise the number of channels we work with: a practical exercise showing this with simulations is in Carucci et al. (2020).

Working with the maps’ native angular resolution. CL23 and other works applied a two-dimensional smoothing to the maps before the cleaning to suppress small-scale systematics. When the smoothing is frequency-dependent, although making all maps share the same resolution, it effectively enhances mode-mixing and eventually degrades the PCA decomposition (Spinelli et al. 2022). We can also smooth all maps with a constant Gaussian kernel that, in this case, would act as a filter, losing cosmological information, too. Sacrificing the small-scale information could be an option for improving the performance of some cleaning methods, although, for instance, Switzer et al. (2013) warned against reducing the ‘angular’ 2D information elements present in the data cube to preserve the efficacy of a contaminant separation based on Eq. (4). Given the already large MeerKAT beam, we did not opt for this choice. We can still filter large and small-scale information with a wavelet kernel and retain them both and even treat them differently within the multiscale framework (i.e. using, at best, the small-scale information and not discarding it: see next point below).

Addressing spatial scales independently. The large scale is where the cosmological signal and the foregrounds are mostly entangled and hard to separate. On the other hand, we see artefacts in the intensity maps on small scales, due to a mixture of systematics and astrophysical foregrounds. Independent handling of contaminant subtraction tasks on different spatial scales with mPCA has proven successful and eventually addresses the long-standing problem in the literature of having to increase the aggressiveness of the cleaning to reduce power on the large scale while typically underestimating the power at the small scale.

9.4. Outlook

We are actively implementing the above best practices in analysing new MeerKAT HI intensity mapping data. However, different cleaning methods might work better with other pre-processing choices than those suitable for the PCA framework. For instance, we know that a generalized morphological component analysis (GMCA, Bobin et al. 2007) can deal with missing channels, leading to discontinuous mixing matrices (Carucci et al. 2020), so additional channel flagging might be beneficial at the cleaning stage¹³. On the other hand, Gaussian process regression (GPR, Mertens et al. 2018) relies mainly on the frequency direction information and its modelling; hence, as in the case of the PCA and SVD (as well as the independent component analysis, ICA), it suffers when channels

¹³ Because GMCA focuses on maximising the sparsity of the \hat{S} components and has no prior on the mixing \hat{A} , the latter matrix can show discontinuous columns, i.e. ‘jumps’ in the spectral structure of the components, to accommodate the lacking frequency information.

are flagged (Soares et al. 2022). Moreover, different strategies have been proposed within the PCA framework to make them more robust against outliers (Zuo et al. 2018) and more flexible, for example by allowing non-linear, functional bases for the decomposition (Irfan & Bull 2021). Room for improvement is present when considering the map resolution, beam deconvolution and beam mismodelling in general that we might want to address before the cleaning; for instance, at the map-making stage (McCallum et al. 2021) or even simultaneously (Carlioni Gertosio & Bobin 2021).

Regarding mPCA and any cleaning performed on the wavelet-filtered data, the choice of wavelet scale (i.e. how we define the large and small scales) is subject to the morphology of the maps, hence survey specific. According to the resolved spatial scale range between the map’s area and the native resolution of the map, one has to test the appropriate choice. For instance, we expect to test and increase the wavelet scales used when extending future survey areas. The choice of the wavelet dictionary, too, is open to discussion (De Caro et al. 2025).

The multiscale approach is generally promising and versatile, and the wavelet filtering delivers an optimal framework for such analysis, namely, the separation of scales facilitates the application of hybrid methods. Here, we choose to keep PCA on both scales analysed, mostly because we want to keep the analysis as simple as possible, while exploring other pre-processing steps; nothing prevents us from mixing different techniques since the two analyses are independent. Other works have proposed ‘hybrid’ and multiscale methods too (Olivari et al. 2016; Spinelli et al. 2022; Podczerwinski & Timbie 2024; Dai & Ma 2025), indeed finding them adaptable to the HI intensity mapping context. In general, a contaminant separation method is more efficient than another if its assumptions suit the data better. Hybrid approaches have the potential to retain the advantages of each of the methods that compose them.

Similarly to PCA, many foreground removal techniques applied in the radio intensity mapping context are based on an eigendecomposition of the data matrix and might benefit from the considerations explored in this article. For example, the Karhunen-Loève (KL) transform is also called the ‘signal-to-noise eigendecomposition’ because it uses priors for the covariances of the cosmological signal and foregrounds (and their ratios) to assign the different eigenvectors to either parts of the observed signal (Bond 1995; Shaw et al. 2014). We find it misleading that sometimes, in the literature, those methods are called ‘linear’ methods as opposed to PCA-like methods which are ‘non-linear’: after all, they all try to isolate the subspace of data where the HI signal overthrows contaminants, with their definition of the filter being different – in the PCA framework the ‘filter’ is unambiguously derived from data through the mixing matrix, that is applied linearly to data to remove contaminants. Similarly to the KL decomposition, methods like GPR are also based on the assumption of knowing at least the ‘shapes’ of the different covariances of the signal and contaminant terms. New adaptations provide for more flexibility (e.g. Diao et al. 2025), and a new machine-learning version of GPR has further relaxed this assumption by allowing the covariances to be learned from simulations (Mertens et al. 2024). Works like the one presented here will help building better simulations and eventually allow us to use simulation-based methods in general.

Finally, we stress again that separation methods such as PCA and mPCA are prone to signal loss. Keeping the loss as low as possible renders results more reliable regardless of how we address it, so it should always be pursued. For instance, in this work, we reached a low enough loss level to be negligible with

our error bars. In perspective, together with pursuing more careful contaminant subtraction, we would need more robust correction approaches or forward-modelling of this effect. For example, we point out that the literature lacks an analysis of the transfer function framework with, for instance, a realistic beam treatment, a crucial aspect that needs attention (some hints of this are in Cunnington et al. 2023b). Indeed, Chen (2025) pointed out that mode-mixing (for which the beam is partly responsible) is the pitfall of the transfer function approach as performed so far. Most importantly, besides the foreground removal process, other pipeline sections (e.g. bandpass calibration and additional filters) can induce signal loss, which must also be estimated, leading to a ‘global’ assessment of the eventual signal loss. At the same time, as we approach an independent clustering detection of the HI intensity maps, we will need to characterise the contribution of residual foregrounds in the cleaned data, which will be competing with signal loss.

10. Conclusions

This work aims to optimise the contaminant suppression scheme for the HI intensity maps from the MeerKAT radio telescope, building on the cross-correlation detection reported in CL23, which serves as our benchmark. We measured cross-spectra not only in agreement with CL23, but also more robust (scale-independent) and with a higher detection significance. Importantly, we successfully addressed and mitigated longstanding issues associated with the PCA framework analysis, specifically, the signal loss and ambiguity in selecting the cleaning level. Overall, this work enhances the confidence in the cosmological signal that we will extract from the upcoming MeerKAT HI intensity mapping data, particularly with respect to preserving the largest scales we can sample.

We have carefully revisited our contaminant subtraction pipeline based on PCA decomposition, reaching a good level of cleaning with a moderate number of removed modes, N_{fg} , thereby resulting in less signal loss. Minor pipeline differences in the HI intensity mapping context can lead to significant differences, as we have experienced first-hand. In this work, we attempted to describe those differences and what changes they bring to the resulting cleaned maps and our interpretations of why this is so.

We have introduced the multiscale mPCA method and are getting encouraging results. Compared to the other pipelines studied, mPCA: (i) delivered the highest amplitude of the cross-correlation (being limited by signal loss, the highest amplitude is the closest-to-truth, Table 1); (ii) is associated with less signal loss when looking at simulations (Fig. 20) and to (iii) less prominent off-diagonal terms when analysing the cleaned cube correlation matrix (Fig. 23) and (iv) lower overall variance (Fig. 22), likely linked to less residual contaminants; (v) is demonstrated to be less sensitive to pre-processing choices (mean-centering and weighting), making it overall more robust; finally, (vi) by splitting the contaminant separation problem into two (at large and small scales), mPCA gives more flexibility in the overall separation process (e.g. Fig. 22). In short, the mPCA approach has revealed improved component separation compared to its whole-scale counterpart PCA, with the prospect of independently detecting a cosmological signal using the intensity map fluctuations.

The focus of this work has been the subtraction of contaminants; hence, we decided to leave everything else unchanged with respect to analysis described in CL23. However, we expect (and we have experienced first-hand) that the data qual-

ity will improve independently of the cleaning because of all the analysis development the collaboration is working on. These works include new calibration schemes and RFI handling, and satellite contamination assessment, as well as new re-gridding schemes and map-making (MeerKlass Collaboration 2025; Cunnington & Wolz 2024; Engelbrecht et al. 2025).

PCA is still the workhorse for the HI intensity mapping cleaning challenge; it is intuitive, computationally cheap, and mathematically sound. Nevertheless, it cannot blindly be applied to data, as its results rely on specific pre-processing choices. The by-products of the cleaning (i.e. the mixing matrix and the components removed) are excellent indicators of the component separation efficiency and their interpretation can guide the modelling of foregrounds and systematics and the development of more sophisticated pipelines, both ‘blind’ and ‘forward-modelling’ alike.

Acknowledgements. The authors would like to acknowledge all members of the MeerKLASS collaboration for the valuable insights and feedback from the ongoing analysis work. IPC would also like to thank Jérôme Bobin for stimulating her interest in component separation and statistical learning. IPC is supported by the European Union within the Next Generation EU programme [PNRR-4-2-1.2 project No. SOE_0000136, RadioGaGa] and acknowledges support in the early stages of this work from the ‘Departments of Excellence 2018–2022’ Grant (L. 232/2016) awarded by the Italian Ministry of University and Research (MUR). JLB acknowledges funding from the Ramón y Cajal Grant RYC2021-033191-I, financed by MCIN/AEI/10.13039/501100011033 and by the European Union “NextGenerationEU”/PRTR, as well as the project UC-LIME (PID2022-140670NA-I00), financed by MCIN/AEI/10.13039/501100011033/FEDER, UE. SC is supported by a UK Research and Innovation Future Leaders Fellowship grant [MR/V026437/1]. MGS acknowledges support from the South African National Research Foundation (Grant No. 84156). JW acknowledges support from the National SKA Program of China (No. 2020SKA0110100). JF acknowledges support of Fundação para a Ciência e a Tecnologia through the Investigador FCT Contract No. 2020.02633.CEECIND/CP1631/CT0002, and the research grants UIDB/04434/2020 and UIDP/04434/2020. YL acknowledges the support of the National Natural Science Foundation of China (grant No. 12473091). AP is a UK Research and Innovation Future Leaders Fellow [grant MR/X005399/1]. We acknowledge the use of the Ilifu cloud computing facility, through the Inter-University Institute for Data Intensive Astronomy (IDIA). The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation.

References

- Abdalla, E., Ferreira, E. G. M., Landim, R. G., et al. 2022, *A&A*, 664, A14
- Ajani, V., Peel, A., Pettorino, V., et al. 2020, *Phys. Rev. D*, 102, 103531
- Alexandris, N., Koutsias, N., & Gupta, S. 2017, *Open Geospatial Data Software Stand.*, 2, 21
- Alonso, D., Ferreira, P. G., & Santos, M. G. 2014, *MNRAS*, 444, 3183
- Amiri, M., Bandura, K., Chen, T., et al. 2023, *ApJ*, 947, 16
- Amiri, M., Bandura, K., Chakraborty, A., et al. 2024, *ApJ*, 963, 23
- 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
- Berti, M., Spinelli, M., Haridasu, B. S., Viel, M., & Silvestri, A. 2022, *JCAP*, 2022, 018
- Bharadwaj, S., Nath, B. B., & Sethi, S. K. 2001, *JApA*, 22, 21
- Blackman, R. B., & Tukey, J. W. 1958, *The Measurement of Power Spectra, from the Point of View of Communications Engineering* (Dover Publications)
- Blake, C., Brough, S., Colless, M., et al. 2010, *MNRAS*, 406, 803
- Blake, C., Brough, S., Colless, M., et al. 2011, *MNRAS*, 415, 2876
- Bobin, J., Starck, J.-L., Fadili, J., & Moudden, Y. 2007, *IEEE Trans. Image Process.*, 16, 2662
- Bond, J. R. 1995, *Phys. Rev. Lett.*, 74, 4369
- Bull, P., Ferreira, P. G., Patel, P., & Santos, M. G. 2015, *ApJ*, 803, 21
- Carloni Gertosio, R., & Bobin, J. 2021, *Digital Signal Process.*, 110, 102946
- Carucci, I. P. 2018, *J. Phys. Conf. Ser.*, 956, 012003
- Carucci, I. P., Irfan, M. O., & Bobin, J. 2020, *MNRAS*, 499, 304
- Chakraborty, A., Datta, A., Roy, N., et al. 2021, *ApJ*, 907, L7
- Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, *Phys. Rev. Lett.*, 100, 091303
- Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, *Nature*, 466, 463
- Chen, Z. 2025, *MNRAS*, 542, L1
- Coles, P., & Jones, B. 1991, *MNRAS*, 248, 1
- Crichton, D., Aich, M., Amara, A., et al. 2022, *J. Astron. Telesc. Instrum. Syst.*, 8, 011019
- Cunnington, S., & Wolz, L. 2024, *MNRAS*, 528, 5586
- Cunnington, S., Irfan, M. O., Carucci, I. P., Pourtsidou, A., & Bobin, J. 2021, *MNRAS*, 504, 208
- Cunnington, S., Li, Y., Santos, M. G., et al. 2023a, *MNRAS*, 518, 6262
- Cunnington, S., Wolz, L., Bull, P., et al. 2023b, *MNRAS*, 523, 2453
- Dai, W.-M., & Ma, Y.-Z. 2025, *ApJS*, 276, 33
- De Caro, B., Carucci, I. P., Camera, S., Remazeilles, M., & Carbone, C. 2025, *JCAP*, submitted [arXiv:2509.02644]
- Diao, K., Grumitt, R. D. P., & Mao, Y. 2025, *ApJ*, 987, 18
- Dickinson, C., Davies, R. D., & Davis, R. J. 2003, *MNRAS*, 341, 369
- Drinkwater, M. J., Jurek, R. J., Blake, C., et al. 2010, *MNRAS*, 401, 1429
- Drinkwater, M. J., Byrne, Z. J., Blake, C., et al. 2018, *MNRAS*, 474, 4151
- Elahi, K. M. A., Bharadwaj, S., Pal, S., et al. 2024, *MNRAS*, 529, 3372
- Engelbrecht, B. N., Santos, M. G., Fonseca, J., et al. 2025, *MNRAS*, 536, 1035
- Favole, G., Granett, B. R., Silva Lafaurie, J., & Sapone, D. 2021, *MNRAS*, 505, 5833
- Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *Astrophys. J.*, 426, 23
- Flöer, L., Winkel, B., & Kerp, J. 2014, *A&A*, 569, A101
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, 433, 181
- Ghosh, A., Bharadwaj, S., Ali, S. S., & Chengalur, J. N. 2011, *MNRAS*, 418, 2584
- Harris, F. 1978, *Proc. IEEE*, 66, 51
- Hartlap, J., Simon, P., & Schneider, P. 2007, *A&A*, 464, 399
- Hu, W., Wang, X., Wu, F., et al. 2020, *MNRAS*, 493, 5854
- Irfan, M. O., & Bobin, J. 2018, *MNRAS*, 474, 5560
- Irfan, M. O., & Bull, P. 2021, *MNRAS*, 508, 3551
- Irfan, M. O., Bull, P., Santos, M. G., et al. 2022, *MNRAS*, 509, 4923
- Irfan, M. O., Li, Y., Santos, M. G., et al. 2024, *MNRAS*, 527, 4717
- Jolliffe, I. T., & Cadima, J. 2016, *Phil. Trans. R. Soc. A.*, 374
- Joseph, R., Courbin, F., & Starck, J. L. 2016, *A&A*, 589, A2
- Kaiser, N. 1987, *MNRAS*, 227, 1
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *Astrophys. J.*, 538, 473
- Li, L.-C., Staveley-Smith, L., & Rhee, J. 2021a, *Res. Astron. Astrophys.*, 21, 030
- Li, Y., Santos, M. G., Grainge, K., Harper, S., & Wang, J. 2021b, *MNRAS*, 501, 4344
- Li, Y., Wang, Y., Deng, F., et al. 2023, *ApJ*, 954, 139
- Loeb, A., & Wyithe, J. S. B. 2008, *Phys. Rev. Lett.*, 100, 161301
- 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, *MNRAS*, 506, 5075
- McCallum, N., Thomas, D. B., Bull, P., & Brown, M. L. 2021, *MNRAS*, 508, 5556
- MeerKlass Collaboration (Barberi-Squarotti, M., et al.) 2025, *MNRAS*, 537, 3632
- Mertens, F. G., Ghosh, A., & Koopmans, L. V. E. 2018, *MNRAS*, 478, 3640
- Mertens, F. G., Bobin, J., & Carucci, I. P. 2024, *MNRAS*, 527, 3517
- Miranda, A. A., Le Borgne, Y.-A., & Bontempi, G. 2008, *Neural Process. Lett.*, 27, 197
- Norberg, P., Baugh, C. M., Gaztañaga, E., & Croton, D. J. 2009, *MNRAS*, 396, 19
- Obuljen, A., Castorina, E., Villaescusa-Navarro, F., & Viel, M. 2018, *JCAP*, 2018, 004
- Offringa, A. R., & Smirnov, O. 2017, *MNRAS*, 471, 301
- Olivari, L. C., Remazeilles, M., & Dickinson, C. 2016, *MNRAS*, 456, 2749
- Pal, S., Elahi, K. M. A., Bharadwaj, S., et al. 2022, *MNRAS*, 516, 2851
- Paul, S., Santos, M. G., Chen, Z., & Wolz, L. 2023, *ApJ*, submitted [arXiv:2301.11943]
- Peel, A., Lin, C.-A., Lanusse, F., et al. 2017, *A&A*, 599, A79
- Pen, U.-L., Staveley-Smith, L., Peterson, J. B., & Chang, T.-C. 2009, *MNRAS*, 394, L6
- Planck Collaboration I. 2020, *A&A*, 641, A1
- Podczewinski, J., & Timbie, P. T. 2024, *MNRAS*, 527, 8382
- Santos, M., Bull, P., Camera, S., et al. 2016, *MeerKAT Science: On the Pathway to the SKA*, 32
- Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014, *ApJ*, 781, 57
- Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2015, *Phys. Rev. D*, 91, 083514
- Shlens, J. 2014, ArXiv e-prints [arXiv:1404.1100]
- Soares, P. S., Watkinson, C. A., Cunnington, S., & Pourtsidou, A. 2022, *MNRAS*, 510, 5872
- Spinelli, M., Carucci, I. P., Cunnington, S., et al. 2022, *MNRAS*, 509, 2048
- Starck, J.-L., Fadili, J., & Murtagh, F. 2007, *IEEE Trans. Image Process.*, 16, 297
- Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *MNRAS*, 434, L46

- Switzer, E. R., Chang, T. C., Masui, K. W., Pen, U. L., & Voytek, T. C. 2015, [ApJ](#), 815, 51
- Tramonte, D., & Ma, Y.-Z. 2020, [MNRAS](#), 498, 5916
- Villaescusa-Navarro, F., Alonso, D., & Viel, M. 2017, [MNRAS](#), 466, 2736
- Wang, J., Santos, M. G., Bull, P., et al. 2021, [MNRAS](#), 505, 3698
- Wang, H., Mena-Parra, J., Chen, T., & Masui, K. 2022, [Phys. Rev. D](#), 106, 043534
- Witzemann, A., Bull, P., Clarkson, C., et al. 2018, [MNRAS](#), 477, L122
- Wolz, L., Blake, C., Abdalla, F. B., et al. 2017, [MNRAS](#), 464, 4938
- Wolz, L., Pourtsidou, A., Masui, K. W., et al. 2022, [MNRAS](#), 510, 3495
- Xu, Y., Wang, X., & Chen, X. 2015, [ApJ](#), 798, 40
- Zuo, S., Chen, X., Ansari, R., & Lu, Y. 2018, [AJ](#), 157, 4

Appendix A: Testing channel flagging and reconvolution

Some pre-processing choices in the present analysis were justified by the evidence in other—simulation-based—works (Sect. 4). Here, with the cross- $P(k)$ measurements, we can test them on data and further support those findings.

Appendix A.1: Channel flagging

The extra channel flagging makes the data cube information discontinuous in frequency. When seen in the context of Eq. (4), this translates into looking for a discontinuous mixing matrix, which the PCA-based methods have difficulty determining. We put this into practice by identifying the channels from the CF data cube that exhibit the highest absolute deviation in variance, the latter fitted in frequency with an order 2 polynomial. We mimic the CL23 analysis and discard ~ 50 channels (maps) from the initial 250. We collapse the data cube and continue the analysis. We show the final results from the cross- $P(k)$ fits as a function of N_{fg} removed in Fig. A.1 with empty triangles to be compared to the ‘standard’ analysis with filled blue circles: results have deteriorated for all methods (PCA in the left panels, PCAw in the middle, and mPCA on the right). The top panels refer to the goodness of the fit, which is underperforming for all cases looked at. Frequency-incomplete data compromise PCA-like contaminant separations. Carucci et al. (2020) show with simulations that the separation does not improve increasing the number of components removed, as we show here for up to $N_{\text{fg}} = 7$.

Appendix A.2: Reconvolution

Reconvolution has been customary in analysing hydrogen intensity mapping experiments (Masui et al. 2013; Wolz et al. 2017; Anderson et al. 2018; Wolz et al. 2022, CL23). It is performed to reduce overall data variance and, specifically, to (i) equalise the cube resolution, since component separation is more effi-

cient when maps share resolution, and (ii) smooth small-scale artefacts. However, (i) a proper beam-deconvolution needs an accurate beam model, but typically, reconvolution assumes a Gaussian kernel; (ii) and it potentially increases mode-mixing by altering maps in a frequency-dependent way, hence detrimental to cleaning efficiency. We test it against our measurements here.

We follow CL23. We convolve each temperature map $T(p, \nu)$ at frequency ν with the following kernel:

$$\mathcal{K}(\Delta p, \nu) = \exp \left[-\frac{\Delta p^2}{2[\gamma\sigma_{\text{max}}^2 - \sigma_{\text{beam}}^2(\nu)]} \right], \quad (\text{A.1})$$

with Δp the angular separation between pixels, and $\sigma_{\text{beam}}(\nu)$ the frequency-dependent resolution when approximating the beam with a Gaussian function as defined in Eq. (16). σ_{max} is the maximum $\sigma_{\text{beam}}(\nu)$ value (corresponding to the lowest frequency in the cube); we set $\gamma = 1.2$ as in CL23, yielding a frequency-independent $\gamma\theta_{\text{FWHM}}(\nu_{\text{min}}) = 1.82$ deg effective resolution, compared to the native median of 1.28 deg we started with. We normalise the kernel in Eq. (A.1) so that the sum over pixels is equal to one. We then perform a weighted reconvolution:

$$T^{\text{reconv}}(p, \nu) = \frac{[T(p, \nu) w_{\text{HI}}(p, \nu)] * \mathcal{K}(\Delta p, \nu)}{w_{\text{HI}}(p, \nu) * \mathcal{K}(\Delta p, \nu)}. \quad (\text{A.2})$$

Finally, the weight maps too are reconvolved according to

$$w_{\text{HI}}^{\text{reconv}}(p, \nu) = \frac{[w_{\text{HI}}(p, \nu) * \mathcal{K}(\Delta p, \nu)]^2}{w_{\text{HI}}(p, \nu) * \mathcal{K}^2(\Delta p, \nu)}. \quad (\text{A.3})$$

After reconvolution, we continued the analysis, going through contaminant separation and measuring the cross- $P(k)$ with the galaxies. Results are displayed as filled green squares in Fig. A.1. Let us start discussing results for the multiscale mPCA analysis in the last column of Fig. A.1. The reconvolution procedure broke down the scales’ separation with the starlet decomposition, the mPCA cleaning failed.

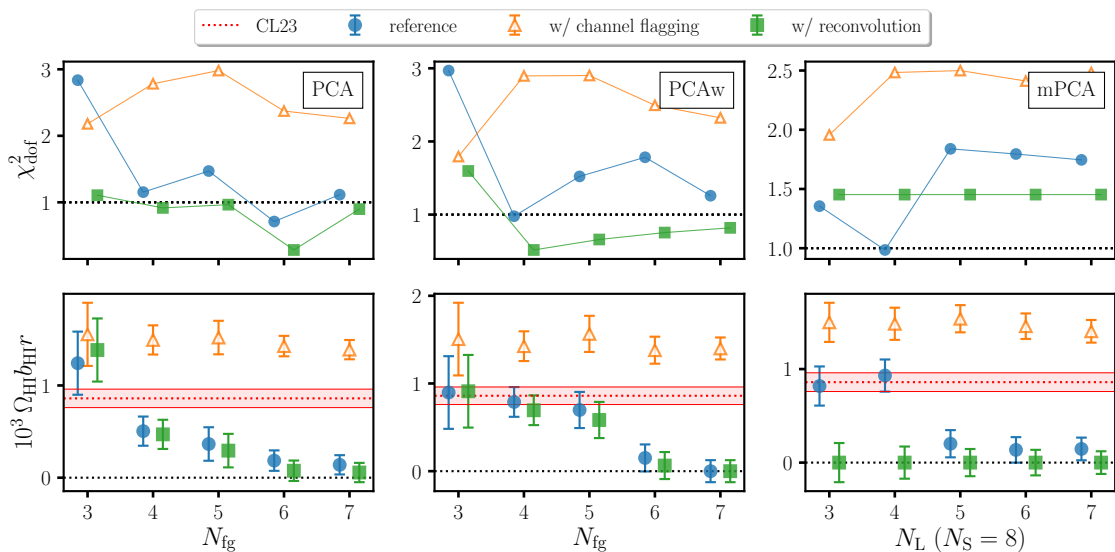


Fig. A.1. Top panels: Reduced χ^2_{dof} of the cross-power best-fits. Bottom panels: Corresponding amplitude values compared to the CL23 result in red, as functions of the cleaning level. Columns correspond to PCA, PCAw and mPCA from left to right. The filled blue circles correspond to the ‘standard’ reference analysis; the empty orange triangle to the analysis with the extra flagging of $\sim 20\%$ of the channels before the component separation; the filled green squares correspond to the analysis performed reconvolving all maps before the component separation.

For PCA and PCAw, the analyses performed after reconvolution are still competitive, although the uncertainties on the measurements increased, χ^2_{dof} are smaller, and the overall amplitudes have also decreased. The frequency-dependent smoothing does not benefit PCA-like contaminant separation.