

H I intensity mapping with the MIGHTEE Survey: first results of the H I power spectrum

Aishrila Mazumder ^{1,★}, Laura Wolz,¹ Zhaoting Chen ^{2,1}, Sourabh Paul ^{1,3,4}, Mario G. Santos ^{3,5},
Matt Jarvis ^{6,3}, Junaid Townsend,³ Srikrishna Sekhar ^{7,8} and Russ Taylor ^{3,8,9}

¹Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

²Institute for Astronomy, The University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK

³Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, Bellville, Cape Town 7535, South Africa

⁴Department of Physics, McGill University, QC H3A 2T8, Montreal, Canada

⁵South African Radio Astronomy Observatory (SARAO), 2 Fir Street, Cape Town, 7925, South Africa

⁶Astrophysics, Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

⁷National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

⁸Inter-university Institute for Data Intensive Astronomy, Department of Astronomy, University of Cape Town, Rondebosch, Cape Town 7701, South Africa

⁹Department of Astronomy, University of Cape Town, Rondebosch, Cape Town 7701, South Africa

Accepted 2025 June 11. Received 2025 June 11; in original form 2024 June 6

ABSTRACT

We present the first results of the H I intensity mapping power spectrum analysis with the MeerKAT International GigaHertz Tiered Extragalactic Exploration (MIGHTEE) survey. We use data covering ~ 4 square degrees in the COSMOS field using a frequency range of 962.5–1008.42 MHz, equivalent to H I emission in $0.4 < z < 0.48$. The data consist of 15 pointings with a total of 94.2 h on-source. We verify the suitability of the MIGHTEE data for H I intensity mapping by testing for residual systematics across frequency, baselines, and pointings. We also vary the window used for H I signal measurements and find no significant improvement using stringent Fourier mode cuts. We compute the H I power spectrum at scales $0.5 \text{ Mpc}^{-1} \lesssim k \lesssim 10 \text{ Mpc}^{-1}$ in autocorrelation as well as cross-correlation between observational scans using power spectrum domain averaging for pointings. We report consistent upper limits of $29.8 \text{ mK}^2 \text{ Mpc}^3$ from the 2σ cross-correlation measurements and $25.82 \text{ mK}^2 \text{ Mpc}^3$ from autocorrelation at $k \sim 2 \text{ Mpc}^{-1}$. The low signal-to-noise ratio in this data potentially limits our ability to identify residual systematics, which will be addressed in the future by incorporating more data in the analysis.

Key words: techniques: interferometric – large-scale structure of Universe – cosmology: observations – radio lines: galaxies.

1 INTRODUCTION

One of the most important objectives of cosmological observations is measuring the distribution and evolution of dark matter in the Universe. This is traditionally done with large galaxy surveys using their positions as tracers of the underlying matter distribution. With the advent of more sensitive instruments, galaxy surveys are being conducted with increased depth and sensitivity to produce tighter constraints on different cosmological parameters (see e.g. Alam et al. 2021; Abbott et al. 2022; Chan et al. 2022). Alternatively, the integrated emission of spectral lines from galaxies can also be used to map the large-scale matter distribution. This method is known as line intensity mapping (IM) and is a very efficient technique for tracing the distribution of matter in the Universe (e.g. Visbal, Trac & Loeb 2011; Kovetz et al. 2017).

Similar to traditional galaxy surveys, IM experiments use biased tracers of dark matter to detect its large-scale distribution. However, unlike galaxy surveys, they do not require the detection of individual

galaxies, but rather the cumulative emission of the target spectral line over large voxels in the sky. Thus, it is faster and can achieve good sensitivity with less observing time than galaxy surveys. IM observations are being done with different lines like carbon monoxide CO (e.g. Cleary et al. 2022), singly ionized carbon C II (e.g. Yue et al. 2015), doubly ionized oxygen O III (e.g. Laporte et al. 2017), and neutral hydrogen H I (e.g. Bharadwaj, Nath & Sethi 2001) for providing constraints on the abundance and clustering of gas in the Universe.

The 21-cm hyperfine transition line from neutral hydrogen (H I) can trace the structure formation out to very high redshifts. Emitted at a rest wavelength of 21.1 cm (or rest frequency of ~ 1420 MHz), it is shifted to lower frequencies due to the expansion of the Universe. This line is detectable using radio telescopes and acts as a valuable cosmological probe at low frequencies. H I constitutes the bulk of the intergalactic medium (IGM) during the early stages of structure formation, i.e. cosmic dawn (CD). When the first stars and galaxies formed, they emitted ionizing radiation ionizing the neutral IGM – a phase transition period called the epoch of reionization (EoR; Furlanetto, Peng Oh & Briggs 2006). In the post-EoR era, the pervading ionizing radiation has kept the IGM ionized until the

* E-mail: aishri0208@gmail.com

present day, with the H I gas only confined inside galaxies and haloes. Thus, we can use this single probe to trace back the history of the formation and evolution of structures in the Universe from the present i.e. redshift $z \sim 0$ to the times the first structures were forming i.e. $z \sim 30$ (Bharadwaj et al. 2001; Battye, Davies & Weller 2004; Furlanetto et al. 2006; Chang et al. 2008; Wyithe, Loeb & Geil 2008; Morales & Wyithe 2010; Spinelli et al. 2020; Chen et al. 2021). Owing to the unique ability of H I to trace different structures (i.e. IGM at high redshifts, dark matter haloes at lower redshifts) depending on the redshift in question, 21-cm observations have become important science drivers for most radio telescope facilities.

H I IM can be done using single dish radio telescopes and radio interferometers. At low redshifts, the former probes the larger angular scales or linear scales ($k \lesssim 0.1 \text{ Mpc}^{-1}$), while the latter probes smaller angular scales or quasi-linear to non-linear scales ($k \gtrsim 0.1 \text{ Mpc}^{-1}$). Forecasts suggest that the SKA Observatory (SKAO), the most sensitive radio telescope to date (currently being built), can make high significance detection of the H I signal between $z \sim 0.5$ and $z \sim 6$ (SKA Cosmology SWG Bacon et al. 2020). However, owing to the inherently weak signal, obtaining a detection is extremely challenging, irrespective of the scales of interest. The deterring factors include many orders of magnitude brighter foregrounds, radio frequency interference (RFI), unmodelled systematics and thermal noise. Foregrounds and systematics contaminate cosmological H I observations irrespective of the redshift targeted. There have been extensive studies on the effect of both foregrounds and systematics for CD/EoR science (e.g. Jelić et al. 2008; Datta, Bowman & Carilli 2010; Chapman et al. 2015; Barry et al. 2016; Trott & Wayth 2016; Ewall-Wice et al. 2017; Mazumder et al. 2022). There have also been studies on how different sources of contamination affect single dish H I IM experiments (for instance, Wolz et al. 2014; Pourtsidou, Bacon & Crittenden 2017; Cunnington et al. 2019; Switzer et al. 2019; Cunnington et al. 2021; Matshawule et al. 2021; Spinelli et al. 2021). Effects of different contaminants on H I IM in the post-EoR era ($z \lesssim 6$) using interferometers have also been done in a few studies (Chen, Wolz & Battye 2022; Chen et al. 2023). Thus, a large amount of dedicated research is ongoing to better understand the contaminants affecting H I IM experiments.

Constraints on the cosmological parameters using single dish telescopes have been produced from the Green Bank Telescope and Parkes. There are detections of power spectrum from H I cross-correlated with optical galaxy surveys at different redshift regimes: $z < 0.1$ (Anderson et al. 2018), $z \sim 0.8$ (Chang et al. 2010; Masui et al. 2013; Switzer et al. 2013), $0.6 < z < 1.0$ (Wolz et al. 2021). The MeerKLASS survey uses the SKAO pathfinder MeerKAT interferometer in single dish mode to access the linear cosmological scales (Santos et al. 2016). This poses many challenges for data processing as demonstrated in Wang et al. (2021) for the pilot data. Nevertheless, cross-correlating the IM pilot data with WiggleZ Dark Energy Survey (Drinkwater et al. 2010) has resulted in the detection of the power spectrum (Cunnington et al. 2022).

There has also been significant progress in performing H I IM with interferometers. While there have been upper limits placed on the H I power spectrum at very high redshift ($7 \lesssim z \lesssim 15$) using telescopes like uGMRT, MWA, HERA, and LOFAR (Paciga et al. 2011; Mertens et al. 2020; Trott et al. 2020; Abdurashidova et al. 2022), no detections have yet been reported. At lower redshifts ($0.78 \lesssim z \lesssim 2.3$), the Canadian Hydrogen Intensity Mapping Experiment (CHIME; Bandura et al. 2014) has detected large-scale structures with H I stacking using Extended Baryon Oscillation Spectroscopic Survey (eBOSS, Dawson et al. 2016) data of galaxies between $0.78 \lesssim z \lesssim 1.43$ (Amiri et al. 2023) and Ly α forest (Amiri et al. 2024) at $z =$

2.3. There are also upper limits on the H I power spectrum using the uGMRT at $1.96 \lesssim z \lesssim 3.58$ (Chakraborty et al. 2021). However, the most compelling result so far has been achieved with the MeerKAT telescope. Detection of the autocorrelation power spectrum has been reported with a ~ 1 square degree deep observation (about 96 h) at two frequency bands centred at redshifts $z \sim 0.32$ and $z \sim 0.44$ (Paul et al. 2023).

In this work, we follow up on the MeerKAT detection using data obtained from the MeerKAT International GigaHertz Tiered Extragalactic Exploration (MIGHTEE) survey of the Cosmic Evolution Survey (COSMOS) field. Covering an area of 4 square degrees using multiple pointings and a total observing time of 94.2 h on source, the data constitutes a small percentage of the full MIGHTEE survey. This data subset is ideal for demonstrating the feasibility of using relatively shallow depth wide area surveys for IM experiments and also provides the first H I power spectrum upper limits from the MIGHTEE data.

The paper is organized in the following manner: in Section 2, we describe the survey and initial data reduction procedure; in Section 3, we describe the analyses and tests performed to check the usability of the MIGHTEE data for H I IM; the results are presented in Section 4; we summarize our findings in Section 5. Throughout this work, we have used the best-fitting cosmological parameters from the Planck 2018 results (Planck Collaboration VI 2020): $\Omega_{\text{M}} = 0.31$, $\Omega_{\Lambda} = 0.68$, $\sigma_8 = 0.811$, $H_0 = 67.36 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

2 MIGHTEE DATA ANALYSIS

The MIGHTEE survey (Jarvis et al. 2016) is an extragalactic radio survey performed with the MeerKAT telescope (Jonas 2018). It covers four extragalactic deep fields – COSMOS, XMM-LSS, E-CDFS, and ELAIS-S1 – using the MeerKAT L band (90–1670 MHz; Goedhart 2020a). The combined area coverage over the four fields is ~ 20 square degree. MIGHTEE survey has several science goals, including spectral line, continuum, and polarization studies. Thus, the fields were chosen to overlap with areas of extensive multiwavelength data. MIGHTEE fields also have both photometric and spectroscopic data, which are detailed in Adams et al. (2020), Bowler et al. (2020), Adams et al. (2021), and Whittam et al. (2023). These fields also coincide with large surveys such as KiDS (de Jong et al. 2013), SDSS (York et al. 2000), GAMA (Driver et al. 2011), as well as newer ones like DES (The Dark Energy Survey Collaboration 2005) and DESI (DESI Collaboration 2022). Given the larger area compared to single-pointing observations and the availability of optical data, the MIGHTEE survey is also suitable for H I IM studies to measure the power spectrum, both in autocorrelation as well as in cross-correlation with optical studies (Paul et al. 2021; Chen et al. 2022). Paul et al. (2021) show that using the full survey of ~ 1000 hours, a detection of the H I signal at $z \lesssim 0.5$ is possible with $\text{SNR} > 7$. However, for this work, we use MIGHTEE data of the COSMOS field for autocorrelation studies.

2.1 Data processing

Fig. 1 shows the layout of the MIGHTEE-COSMOS field, with the black circles extending to the primary beam width of the MeerKAT L-band and the phase centres of each of the 15 pointings denoted by the blue dots. The total observing time per pointing is 8 h, with 6.2 h on-source and a time resolution of 8 s (see Table 1). Hence, there is 120 h of total observation time, with 94.2 h on-source. The pointing centres are arranged in a close-packed mosaic to coincide with the region of




Figure 1. Schematic diagram of the close-packed mosaic observation of MIGHTEE-COSMOS. The black circles represent the FWHM of the MeerKAT *L* band primary beam ($\sim 1^\circ$), with the coloured dots indicating the phase centres of the pointings. The pointing centres are labelled according to their respective COSMOS numbers. Three pointings share a common declination with five different declinations used as phase centres for the MIGHTEE-COSMOS data.

the COSMOS field with the best multiwavelength coverage. For this work, we use the survey data with a frequency resolution of 26 kHz across 32 768 channels in the MeerKAT *L* band, primarily used for MIGHTEE-H I galaxy science (Heywood et al. 2024). Early science data from the MIGHTEE survey used 4096 channels (Heywood et al. 2021).

The COSMOS pointings cover declinations between $+01^\circ 51' 08''.2$ and $+02^\circ 33' 33''.8$. The detailed observation parameters are tabulated in Table 1. Each group of three pointings have their phase centre at the same declination (specified in the last column of Table 1), and for each declination group, the right ascensions are arranged in decreasing order. Throughout the paper, the pointings are arranged identically to Fig. 1. COSMOS.2 has slightly higher on-source integration than the other pointings; we will use COSMOS.2 for discussing results for a single pointing.

The observations used here consist of interleaved target scans of around an hour between calibrator observations for each pointing, giving seven scans per pointing. The combined observing time of the MIGHTEE-COSMOS field is similar to the observing time for the first detection of H I power spectrum using MeerKAT observation of the DEEP2 field (Paul et al. 2023). However, DEEP2 consisted of a single deep pointing instead of multiple pointings used here. The total area covered by this close-packed mosaicing strategy is ~ 4 square degrees.

The data processing steps are described in detail in Heywood et al. (2024) and are briefly outlined here for completeness. The data has a high native frequency resolution of 26 kHz, and is primarily observed for MIGHTEE-H I galaxy science. The target visibilities are Level-1 (L1) calibrated visibilities recovered from KAT Data Access Library (KATDAL¹). These L1 visibilities have reference

calibration solutions applied to the targets using the SARAO Science Data Processor (SDP; Ratcliffe 2020). The bandpass solutions are derived from the primary calibrator (J0408–6545). The time-dependent complex gain solutions are from the secondary calibrator (J1008+0740). The same calibrator sources are used across all pointings.

The calibration report for the L1 visibilities for each pointing can be obtained from the SARAO archive² using the observation IDs listed in Table 1. We summarize the important details here. The initial flags remove known RFI-contaminated channels (see Goedhart 2020b for details). Then, the calibration pipeline performs further flagging to remove time-steps and baselines affected by RFI. For the channels of our interest (i.e. between 962.5 to 1008.4 MHz), the RFI flagging step usually results in a flagging percentage $\lesssim 25$ per cent. The subsequent calibration steps use antenna m060 as the reference antenna for deriving solutions. The delay and bandpass calibration is performed in the frequency range 973–1000 MHz using the primary calibrator (J0408–6545). Inspection of the solutions shows them to be stable and smooth for the COSMOS pointings. J1008+0740 i.e. the secondary calibrator is used to derive the time-dependent complex gain solutions. These solutions were applied to the targets i.e. the COSMOS fields and retrieved these calibrated data as measurement sets. This data is subsequently split into two subbands – LOW (960–1150 MHz, channels averaged to 104.5 kHz resolution) and MID (1290–1520 MHz, channels kept at native resolution of 26.1 kHz).

These calibrated visibilities are further processed in the following steps:

(i) First round of flagging using TRICOLOUR package (Hugo et al. 2022). It uses the SUMTHRESHOLD algorithm (Offringa et al. 2010) to clip values where the sum within a time step and channel exceeds a threshold calculated iteratively. This round flags residual low-level features.

(ii) Construction of a deconvolution mask from the deep continuum image as a part of MIGHTEE continuum science (Hale et al. 2025; Heywood et al. 2024).

(iii) Image plane deconvolution using WSCLEAN to create an image of size 10240×10240 and pixel size 1.1 arcsec with robust weighting ($r = -0.3$).

(iv) CLEAN component model obtained from the previous step is interpolated in frequency on to a higher number of subbands (909 for LOW band used here) to create a spectrally smooth model.

(v) The CLEAN components of the frequency-dependent model images are used in the WSCLEAN predict mode to create model visibilities for each subband.

(vi) One round of phase and delay self-calibration with 32 s time interval and entire bandwidth for each of the MID and LOW subbands as frequency interval.

(vii) Visibility domain continuum subtraction by applying the best-fitting gain solutions from the self-calibration step to the data and subtracting the model visibilities.

(viii) Another round of flagging on residual data.

Each pointing of the COSMOS data undergoes this process to produce the final calibrated visibilities used here. The MIGHTEE H I galaxy detection analysis involves additional imaging steps described in Heywood et al. (2024). However, our approach to IM analysis does not require imaging. Throughout the rest of the paper, we use these

¹<https://github.com/ska-sa/katdal>

²<https://archive.srao.ac.za>

Table 1. Observation details of the MIGHTEE COSMOS Pointings used for this work, arranged in descending order of declination. The constant declination pointings are grouped together with the common declination specified in the last column.

ID	Field	Pointing centre	Date	Time (UTC)	On-source (h)	Declination (deg)
1 585 928 757	COSMOS_9	(10 ^h 01 ^m 54 ^s , +02 ^d 33 ^m 33 ^s .79)	30.04.2020	16:07:25.9–23:36:02.6	6.25	
1 585 844 155	COSMOS_8	(10 ^h 00 ^m 29 ^s , +02 ^d 33 ^m 33 ^s .79)	02.04.2020	16:37:27.3–00:06:03.9	6.25	2.56
1 586 016 787	COSMOS_10	(09 ^h 59 ^m 04 ^s , +02 ^d 33 ^m 33 ^s .79)	04.04.2020	16:35:26.7–00:04:19.4	6.25	
1 621 083 675	COSMOS_4	(10 ^h 01 ^m 11 ^s .063, +02 ^d 22 ^m 57 ^s .39)	15.05.2021	13:25:05.08–20:53:58.4	6.25	
1 619 963 656	COSMOS_2	(09 ^h 59 ^m 46 ^s .15, +02 ^d 22 ^m 57 ^s .39)	02.05.2021	14:17:24.7–21:46:17.3	6.40	2.38
1 586 188 138	COSMOS_11	(09 ^h 58 ^m 21 ^s , +02 ^d 22 ^m 57 ^s .39)	06.04.2020	16:10:28.0–23:39:04.6	6.25	
1 585 498 873	COSMOS_6	(10 ^h 01 ^m 54 ^s , +02 ^d 12 ^m 20 ^s .99)	29.03.2020	16:43:19.7–00:11:48.4	6.25	
1 587 911 796	COSMOS_0 [†]	(10 ^h 00 ^m 28 ^s .6, +02 ^d 12 ^m 21 ^s)	26.04.2020	14:58:02.3–22:26:38.9	6.25	2.21
1 585 413 022	COSMOS_5	(09 ^h 59 ^m 04 ^s , +02 ^d 12 ^m 20 ^s .99)	28.03.2020	16:52:28.9–00:21:13.5	6.25	
1 622 376 680	COSMOS_3	(10 ^h 01 ^m 11 ^s .053, +02 ^d 01 ^m 44 ^s .60)	30.05.2021	12:34:21.9–20:03:22.5	6.25	
1 617 809 470	COSMOS_1	(09 ^h 59 ^m 46 ^s .15, +02 ^d 01 ^m 44 ^s .60)	07.04.2021	15:53:37.9–23:22:22.5	6.25	2.03
1 586 274 966	COSMOS_12	(09 ^h 58 ^m 21 ^s , +02 ^d 01 ^m 44 ^s .60)	07.04.2020	16:17:33.1–23:46:17.7	6.25	
1 586 791 316	COSMOS_14	(10 ^h 01 ^m 53 ^s , +01 ^d 51 ^m 08 ^s .20)	13.04.2020	15:44:06.4–23:12:43.0	6.25	
1 585 671 638	COSMOS_7	(10 ^h 00 ^m 29 ^s , +01 ^d 51 ^m 08 ^s .20)	31.03.2020	16:42:45.4–00:11:30.0	6.25	1.85
1 586 705 155	COSMOS_13	(09 ^h 59 ^m 04 ^s , +01 ^d 51 ^m 08 ^s .20)	12.04.2020	15:48:05.2–23:16:57.8	6.25	

[†]This pointing is originally labelled COSMOS, but has been relabelled here as COSMOS.0 for consistency.

continuum-subtracted visibilities. We are interested in the redshift range of $0.2 \lesssim z \lesssim 0.5$, thus we use the LOW band for this work.

2.2 Visibility domain gridding

The MIGHTEE-COSMOS data consists of continuous target scans of about an hour each, totalling ~ 6 h on-source tracks per pointing. For each pointing, we grid the uv -plane into discrete cells to generate gridded visibility cubes in the (u, v, ν) , i.e. baseline-frequency domain. Each grid has size $\Delta u = \Delta v = 60\lambda$ (where λ corresponds to the central wavelength). The gridding is performed as a function of frequency, with each visibility assigned a single (u, v, ν) cell. This gridding approach can change the (u, v) cell assigned to each visibility as a function of frequency. The visibilities include only the baselines with $\lesssim 20$ per cent flagged data. The chosen cell size (calculated at the band centre) makes the primary beam width in k_{\perp} space negligible, thereby minimizing the effect of mixing different k_{\perp} modes. Each scan, i.e. interleaved time steps between calibrator observations, is gridded separately for each pointing. Fig. 2 shows the distribution of the visibilities in the gridded uv plane for a particular pointing (COSMOS_2) for the full on-source observation. It is evident from Fig. 2 that the uv -cells are not sampled with a uniform density far away from the centre of the plane. MeerKAT has more short baselines than long ones. Thus, across an observing session, the sampling is sparser as one goes away from the centre of the uv plane. The observed fields being located at declinations that remain above the horizon for less time than those further south³ also contribute to the sparse sampling. Additionally, the individual sources are resolved at long baselines, making the Poisson sampling noise dominant over the clustering signal. All these effects combined justify excluding long baselines for the HI power spectrum measurements.

Grid points shown by the black dashed circle in Fig. 2 fall within $\pm 1000\lambda$ from the centre of the uv -plane (corresponding to $k_{\perp} = 10 \text{ Mpc}^{-1}$). This region is uniformly dense for all the pointings. Thus, using only visibilities within a uv -distance of 1000λ results in a uniform sampling of k_{\perp} modes in the final power spectrum estimates. The various analyses performed in the power spectrum




Figure 2. Gridded uv plane for the full on-source observation track of COSMOS_2. Each cell has a size $u = v = 60\lambda$. The black dashed circle is the zone between $\pm 1000\lambda$, used for the PS analysis.

domain throughout the paper utilize the gridded visibilities described here.

2.3 Power spectrum estimator

In this work, we use the calibrated MIGHTEE-COSMOS data to estimate the HI power spectrum. Each pointing has a different phase centre with a relatively low signal-to-noise ratio, and we average over all the pointings, which reduces the noise to enable a combined HI power spectrum measurement. In this analysis, we incoherently average the power spectra measurement in Fourier space. This is different from coherent averaging in the visibility domain done for deep integration over a single pointing such as the DEEP2 analysis (Paul et al. 2023). In this section, we describe its formalism. For the remainder of the article, unless stated explicitly, averaged or

³For details on track length as a function of declination (see Goedhart 2020a).

combined power spectrum refers to the power spectrum estimated from incoherently averaged data.

To calculate the power spectrum, the calibrated interferometric visibilities $V(b, t, \nu)$ are gridded as described in Section 2.2. We then Fourier transform the gridded cube along the frequency axis, resulting in gridded visibility in the Fourier domain, i.e. $\tilde{V}(b, t, \eta)$. This process is similar to the ‘delay spectrum’ approach (Morales & Hewitt 2004; Parsons et al. 2012) of calculating the power spectrum from the visibilities. The gridded visibility cube linearly combines data from multiple baselines in the power spectrum estimation process (a similar method is also used for imaging before Fourier transforming visibilities). Thus, following the convention in Morales et al. (2018), we estimate the ‘reconstructed’ power spectrum in this work.

The Fourier conjugate of the frequency, η is related to k_{\parallel} , i.e. Fourier modes along the line-of-sight ratio via $k_{\parallel} = \frac{2\pi\eta H(z)}{\lambda_{21}(1+z)^2}$, where λ_{21} is the rest-frame wavelength of the 21-cm spin-flip transition line, $H(z)$ is the Hubble parameter corresponding to the redshift z at the central observing frequency ν .⁴ The angular Fourier modes (perpendicular to the line of sight) k_{\perp} is related to the baseline vector b as $k_{\perp} = \frac{2\pi b}{\lambda X}$, where λ is the wavelength corresponding to ν and X is the co-moving distance to the corresponding redshift.

The observed visibilities for a patch of sky with the same \mathbf{k}_{\perp} at a certain frequency should have the same value (plus a thermal noise component) at any given time. Hence, in the Fourier domain, $\tilde{V}(b, t, \eta)$ can be written as $\tilde{V}(k)$, with $k = (k_{\perp}, k_{\parallel})$. The corresponding 3D power spectrum P_D is given by Morales & Hewitt (2004) and Parsons et al. (2012):

$$P_D(k_{\perp}, k_{\parallel}) = \left(\frac{X^2 Y}{\Omega_{ps} B} \right) \left(\frac{\lambda^2}{2k_B} \right)^2 \text{Re}\{\tilde{V}_i(k_{\perp}, k_{\parallel}) \tilde{V}_j^*(k_{\perp}, k_{\parallel})\}, \quad (1)$$

where k_B is the Boltzmann constant, Y represents the comoving depth along the line-of-sight ratio corresponding to bandwidth B . Ω_{ps} is the power squared primary beam, defined as $\Omega_{ps} = \int d\ell dm |A(\ell, m)|^2$, with $A(\ell, m)$ being the primary beam pattern, as opposed to the standard integrated power of the beam obtained by an integration over the beam i.e. $\int d\ell dm A(\ell, m)$ (see Parsons et al. 2014 for details).

For each pointing, we can estimate the power spectrum either as autocorrelation or as cross-correlation between different observational scans. For the autocorrelation, we compute the gridded visibilities via $\tilde{V}_j = \tilde{V}_i^*$ in equation (1). The autocorrelation power spectrum contains an additive thermal noise bias term along with the observed sky signal. However, we can infer the noise in the data using observation parameters and remove it from the autopower spectrum to estimate the noise-free signal power spectrum. In the case of cross-correlation, we create two visibility sets from the observational scans. For example, we can create \tilde{V}_i and \tilde{V}_j from alternate 1-h scans. They are split into ‘even scans’ and ‘odd scans’ based on the parity of the respective scan numbers, i.e. scan numbered 0 is even, scan numbered 1 is odd and so on. The cross-correlation removes the uncorrelated noise bias from the independent samples, in principle measuring the signal power spectrum. We conducted tests on the impact of the time-scales for creating the cross-correlation samples, and we could not identify significant differences in the cross-correlation amplitudes within the limited signal-to-noise ratio of our observations. The 3D powers from both auto and cross-correlations are binned into cylindrical or spherical bins to give the corresponding power spectra. Each pointing of MIGHTEE-COSMOS observes a slightly different

part of the sky, so the 3D powers obtained in equation (1) are averaged over all the pointings to get the final averaged power spectrum.

For incoherent averaging, we follow the steps below:

- (i) Fourier transform the gridded visibilities along the frequency axis per pointing to obtain the 3D power spectrum.
- (ii) Weight the 3D k cell by the respective baseline density for each pointing. The weighting accounts for the differences in baseline coverage among the pointings.
- (iii) Average all weighted 3D power spectra to obtain the averaged 3D power.
- (iv) Bin the averaged 3D power spectrum into spherical/cylindrical bins.

The final combined 3D power spectrum is

$$P_D(k_{\perp}, k_{\parallel}) = \frac{\sum_{i=1}^{15} w_i(k_{\perp}, k_{\parallel}) P_D^i(k_{\perp}, k_{\parallel})}{\sum_{i=1}^{15} w_i(k_{\perp}, k_{\parallel})}, \quad (2)$$

where the i is over all 15 pointings and $P_D^i(k_{\perp}, k_{\parallel})$ is the 3D power for the i th pointing, weighted by the factor w_i which is the baseline density per 3D k pixel.

The 2D cylindrical power spectrum, $P(k_{\perp}, k_{\parallel})$ is calculated for the MIGHTEE COSMOS data binning the averaged 3D power spectrum, $P_D(k_{\perp}, k_{\parallel})$ into cylindrical bins in the line-of-sight and transverse Fourier modes, weighted by the inverse noise variance giving the estimator

$$P^i(k_{\perp}, k_{\parallel}) = \frac{\left(\sum_j P_D^j(k_{\perp}, k_{\parallel}) s_j^{iv}(k_{\perp}, k_{\parallel}) \right)}{\left(\sum_j s_j^{iv}(k_{\perp}, k_{\parallel}) \right)}, \quad (3)$$

where the weight s_j^{iv} is the inverse noise variance $1/\sigma_j^2(k_{\perp}, k_{\parallel})$, j loops over all the k modes in the i th k bin. Owing to the differences in observing days and pointing centres, we simulate the noise variance for each pointing separately (discussed in Section 3.2). The 3D power spectrum is spherically averaged to obtain the statistically relevant 1D or spherical power spectrum using Fourier modes outside the foreground-dominated region or the so-called foreground wedge (Datta et al. 2010).

2.3.1 Foreground avoidance

For a sky with pure statistically isotropic and homogeneous HI signal, the power spectrum should be a function of $k = \sqrt{k_{\perp}^2 + k_{\parallel}^2}$. However, real observations will have contributions from systematics and thermal noise (which is highly dependent on the baseline distribution for the particular observation). The isotropy of HI signal is also broken by redshift space distortions. This effect reduces the amplitude of fluctuations at large k_{\parallel} at the small spatial scales observed by an interferometer. Astrophysical foregrounds are spectrally smooth, their interaction with the instrument chromaticity concentrates power at small k_{\parallel} or short delays creating the foreground wedge. In interferometric observations, the wedge allows us to isolate the foregrounds from the HI signal using foreground avoidance (Datta et al. 2010; Morales et al. 2012; Parsons et al. 2012; Trott, Wayth & Tingay 2012; Vedantham, Shankar & Subrahmanyan 2012; Abdurashidova et al. 2022; Paul et al. 2023). Foreground avoidance, in essence, excludes the foreground dominated (k_{\perp}, k_{\parallel}) modes inside the wedge for calculating the spherical power spectrum. This method utilizes data in the foreground-free region and does not require

⁴The redshift z is given by $1+z = \nu_{21}/\nu$

additional foreground removal. This work uses foreground avoidance to measure the final spherically averaged power spectrum $P(k)$.

For an instrument with a primary beam size θ_B , foregrounds should ideally be confined within the ‘horizon limit’ (Liu, Parsons & Trott 2014), given by

$$k_{\parallel} = \frac{XH(z)\sin\theta_B}{c(1+z)}k_{\perp}, \quad (4)$$

where $H(z)$ is the Hubble parameter. Using equation (4) for calculating the primary beam limit, we get $k_{\parallel} \sim 0.01k_{\perp}$ for the MeerKAT frequency band used here. However, due to imperfections in the instrument, the presence of calibration errors and leakage from beyond the instrument beam, the contamination can extend beyond these modes and start affecting higher modes as well (Datta et al. 2010; Mazumder et al. 2022). Thus, we employ the conservative horizon limit, i.e. $\sin\theta = 1$, giving $k_{\parallel} \sim 0.3k_{\perp}$. We also use different k mode cuts (discussed in detail in Section 3.3.1) to test for improvements in our final measurements by excluding more potentially contaminated modes.

The uncertainties in the calculated power for each k mode of the final power spectrum are calculated from the variance of the power in each bin as

$$(\Delta P_D^i(k_{\perp}, k_{\parallel})) = \sqrt{\frac{\left(\sum_j (P_D^j(k_{\perp}, k_{\parallel}) - P_D^i(k_{\perp}, k_{\parallel}))^2 s_j^2(k_{\perp}, k_{\parallel})\right)}{\left(\sum_j s_j(k_{\perp}, k_{\parallel})\right)^2}}, \quad (5)$$

where the index j loops over all modes falling in the i th k -bin.

2.3.2 Noise in incoherent averaging

As described above, the power spectrum measured with incoherent averaging is the so-called square-then-average method as opposed to the coherent averaging or the average-then-square method (Liu, Zhang & Parsons 2016). The noise scales down slower in the former method compared to the latter. Roughly, for N pointings, the noise amplitude scales down by $1/\sqrt{N}$, since N independent measurements are averaged after squaring (as opposed to averaging N independent realizations and then squaring, which reduces the noise by $1/N$). Thus, for ~ 100 h of data over 15 pointings, the noise is equivalent to ~ 25 h or ~ 4 times that of a single 100-h pointing. Thus, despite an on-source time of 94.2 h for MIGHTEE-COSMOS (similar to ~ 96 h for DEEP2 used in Paul et al. 2023), the expected signal-to-noise ratio is that of ~ 24 h of integration. Using the radiometer equation, for the used bandwidth of 46 MHz, we expect a root-mean-squared thermal noise fluctuation of ~ 5 mK for a single pointing of 94.2 hs and about 10 mK if we divide the time into 15 pointings. Simulations from Paul et al. (2021) show that using the full ~ 1000 h of the MIGHTEE survey, H I IM power spectrum detection can be obtained even by incoherent averaging. In this work, we use a part of the full survey data and obtain the first upper limits from MIGHTEE observations.

3 QUALITY ASSESSMENT FOR MIGHTEE DATA

The redshifted 21-cm signal is an inherently weak signal and prone to contamination by several factors. The data calibration, as described above, invariably leaves some residual systematics and possibly, low-level RFI features. The MIGHTEE-COSMOS pointings are

spread over a large area (~ 4 square degree), with possibilities for declination-dependent systematic contamination. In this section, we present the data quality assessments to check for such effects for individual pointings and the combined data.

3.1 Residual RFI

MeerKAT L -band data has several frequency channels pre-flagged by SARAO since they are affected heavily by known RFI contamination (Goedhart 2020b). The extra flagging steps described in Section 2.1 remove additional RFI contaminants. As a first check for data quality, the total flagging percentage for the entire LOW band as a per pointing is shown by blue points in Fig. 3. We see that the percentage of flagged data is below 30 per cent for all pointings using the full LOW band, with the highest flagging percentage in COSMOS_3 (~ 27 per cent). The overall trend across most of the pointings shows $\lesssim 15$ per cent flagging, with the exceptions COSMOS_1, COSMOS_2, COSMOS_3, COSMOS_4, and COSMOS_13. Since there are multiple pointings at the same declination, we checked for any declination-dependent variation in the flagging percentage. There is no indication of such effects from the flagging percentages since COSMOS_2, COSMOS_4, and COSMOS_11 are at 2.38° , but the latter has ~ 8 per cent data flagged. Similarly, COSMOS_1 and COSMOS_3 show higher flagging percentages than COSMOS_12 (~ 6 per cent); but all three are at a declination of 2.03° . The same is also true for COSMOS_13 when compared to COSMOS_7 and COSMOS_14 at 1.85° .

For further investigation, we plot the flagging percentage as a function of frequency for each pointing. Fig. 4 shows the percentage of data flagged per channel of the LOW band for the pointings as indicated in each panel. Constant declination pointings are grouped in rows. In Paul et al. (2023), the cleanest part of the band was obtained between 962.55–1008.42 MHz (centred at ~ 986 MHz) for redshift ~ 0.44 and between 1054 and 1100 MHz, i.e. centred at ~ 1077.5 MHz for $z \sim 0.32$. The respective regions are marked by dashed and dotted lines in Fig. 4. The higher frequency subband

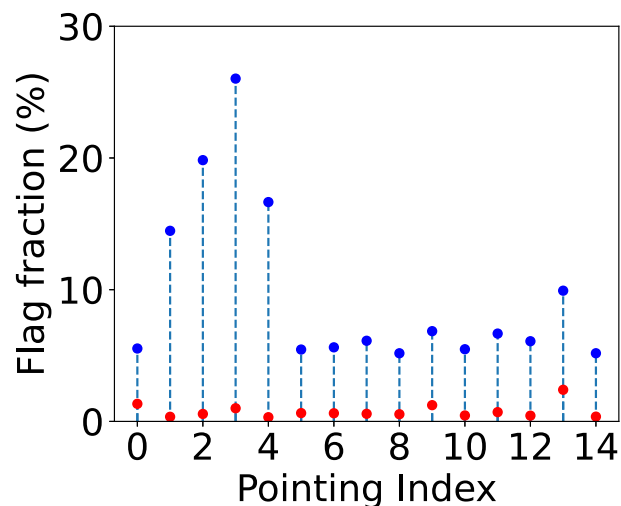


Figure 3. Flagging fraction of individual pointings over the entire LOW band (blue points). The overall flagging percentage is < 30 per cent for all 15 pointings. The red points are plotted for a smaller subband corresponding to a frequency range 962.5–1008.4 MHz, with an additional criterion of excluding baselines with more than 20 per cent flagging. This results in < 10 per cent flagging for all cases.




Figure 4. Flagging fraction of individual pointings as a function of frequency of the full LOW band. The region enclosed by black dashed lines is the frequency range for $z \sim 0.44$, and the black dotted line is that for $z \sim 0.32$ used in Paul et al. (2023). Panels in each row represent pointings at the same declination. The columns are arranged from left to right in order of right ascensions as shown in Fig. 1 and rows are in descending order of declination.

is reasonably clean for all the MIGHTEE pointings, but the lower subband shows large flagging fractions in some cases, e.g. for pointings COSMOS.1, COSMOS.2, COSMOS.3, and COSMOS.4. Ambient factors like the presence of the sun or observations taken during local sunset (or sunrise) can result in additional features leading to excess flagging. However, many of the pointings have the sun present, while COSMOS.13 and COSMOS.14 also cover local sunset but do not show such high amounts of flagging. Further investigations show that the four pointings were taken a year after the others, hinting at the presence of some excess ambient interference leading to more flagging for these specific observations.

Flagging percentages can also vary over baselines for the same observation block. Since a very large amount of flagged data may produce spurious features in the final power spectrum, we filter out baselines having more than 20 per cent flagged data for our analyses. Red points in Fig. 3 show the flagging percentage for each pointing in the dashed subband (i.e. centred at ~ 986 MHz) using only baselines

with ≥ 80 per cent unflagged data. The flagging percentage as a function of channels per pointing is plotted in Fig. 5 for the LOW band. The Y-axes are zoomed-in to 10 per cent for better visualization. The subband corresponding to $z \sim 0.32$ (1054–1100 MHz) is shown by dotted boundaries. It is seen that while excluding heavily flagged baselines does improve the overall behaviour, this subband still has a few channels with flagging fractions higher than 10 per cent, specifically towards the higher frequency end. Conversely, the shaded region in Fig. 5 (with frequency 962.55–1008.42 MHz corresponding to $z \sim 0.44$) shows that imposing the baseline cut makes the subband cleaner while maintaining an almost uniform flag level throughout for almost all the pointings. Thus, we confine our subsequent analysis for this work to $z \sim 0.44$ subband i.e. the shaded regions in Fig. 5, which is consistent with one of the subbands used in Paul et al. (2023). Unlike Paul et al. (2023), we have not performed any in-painting for the flagged data, since, as mentioned, the flagging percentages are almost uniform across the subband. Future work will investigate other frequency windows and how the power spectrum is affected




Figure 5. Flagging fraction of individual pointings as a function of the frequency of the LOW band excluding all baselines which are ≥ 20 per cent flagged. The Y-axes are limited to 10 per cent for better visualization. The region enclosed by the black dotted circle is $z \sim 0.32$ which shows > 10 per cent flagging for most pointings. The shaded region is for $z \sim 0.44$ which shows $\lesssim 10$ per cent flagging for most pointings, thus used for further analysis. Panels are arranged the same as Fig. 4.

by incorporating visibilities with higher flagging levels and data inpainting.

3.2 Thermal noise simulations

This section describes the methodology to simulate the noise of our observations. The noise estimation is important because the power spectrum estimator described in Section 2.3 uses inverse noise variance weighting. Since astrophysical foregrounds in our scales of interest are expected to have negligible circular polarization, we can use Stokes V data to estimate the thermal noise amplitude. It was shown in Paul et al. (2021, 2023) that the MeerKAT data has thermal noise consistent with the Stokes V. Gridding the Stokes V data similar to the total intensity data as described in Section 2.2, the noise amplitude for N baselines in a uv-cell with average visibility V is $\sigma_N = \text{std}(V\sqrt{N})$.

To estimate the thermal noise per pointing, we simulate the noise power spectrum where we replace the gridded visibilities in each uv-cell (as described in Section 2.2) with a random realization of a zero mean complex circular Gaussian with a standard deviation of

$$\sigma_N = \frac{2k_B T_{\text{sys}}}{A_e \sqrt{\delta\nu \delta t}}, \quad (6)$$

where k_B is the Boltzmann constant, the channel width $\delta\nu$ is 104.5 kHz and time resolution is $\delta t = 8$ s. Since we are generating the noise from the gridded visibilities, the generated noise realizations are divided by the square root of the counts per cell. The natural sensitivity of the instrument is A_e/T_{sys} with system temperature T_{sys} and effective area A_e having anticipated values of $6.22 \text{ m}^2 \text{ K}^{-1}$ for MeerKAT (Goedhart 2020a). The expected noise is also calculated from Stokes V as described above, and we obtain a value of $\sim 6.49 \text{ m}^2 \text{ K}^{-1}$.

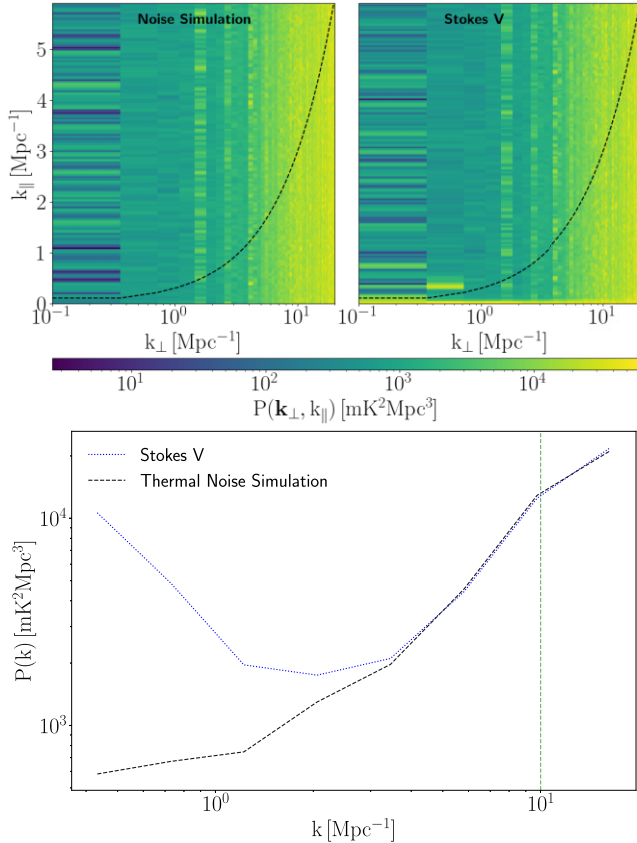


Figure 6. Top: Cylindrical power spectrum for a single realization of thermal noise simulation (left panel) and gridded Stokes V visibilities (right). Bottom: The spherically averaged power spectrum of the thermal noise simulation compared against Stokes V. The green dashed curve is the k cut-off used for final power spectrum measurements. The simulated noise is consistent with Stokes V, except in the lower k_\perp – k_\parallel plane, as evident from both panels.

We simulate 10^5 noise realizations for each pointing and the average value for each k pixel gives the noise variance. Fig. 6 shows the Stokes V and noise power spectra for COSMOS_2 with the cylindrical power on the top panel and the spherical power on the bottom panel. The top-left panel demonstrates the noise simulations, top-right shows the Stokes V power on the right. The black dashed curve on the bottom panel shows the spherical power spectrum of the noise simulations while the blue dotted curve shows the same for Stokes V. From the top panel, the Stokes V power closely resembles the simulated noise power except at the lowest region on the k_\perp – k_\parallel plane. This is also reflected in the bottom panel of Fig. 6, where the simulated spherical noise power deviates from Stokes V at small scales. The value of k for which Stokes V deviates from simulated thermal noise estimates is similar to that seen in Paul et al. 2023. Although Fig. 6 shows a particular pointing, the same trend is seen for all the pointings of COSMOS data.

The sparser sampling of long baselines (or large k_\perp values) compared to the short ones leads to higher noise in the former case. However, the line of sight i.e. k_\parallel direction does not have this non-uniformity. Thus, the noise is expected to vary more in the k_\perp direction, which predominantly contributes to the noise power spectrum. This leads to the increase in both the noise power and Stokes V powers seen at large k values. The green dashed line at $k \sim 10$ Mpc $^{-1}$ corresponds to k_\perp beyond which baselines are sparsely

sampled. Any measurement beyond these values would be noise-like, demonstrating the need for a cut-off value for the highest usable k_\perp value for our final power spectrum measurements.

The deviation between the simulated noise and Stokes V is due to power leakage into Stokes V. The structure is clearly visible at low k_\perp , as seen in both Figs 6 and A1. These modes are dominated by contributions from the central part of the uv plane, resulting in high signal-to-noise ratio and more evident systematic features. Conversely, the large k values being noise-dominated, we do not see this feature so clearly.

Fig. A1 shows the ratio of the Stokes V power over the noise power for all MIGHTEE-COSMOS fields, all of which show the excess at the smallest k_\parallel across all k_\perp . However, we use only the longer baselines (where the noise simulations and Stokes V match quite well) to normalize the simulated noise power. Thus, this is not a cause of concern for our results.

3.3 Power spectrum analysis

In this section, we present tests performed in the power spectrum domain. We start with the 3D power spectrum (equation (1)) obtained from gridded continuum-subtracted visibility cubes as described in Section 2.2. The power spectrum is multiplied by a Blackman–Harris window, which suppresses the foreground leakage at higher k_\parallel modes. We use both auto and cross-power spectra to test for residual systematic contributions.

3.3.1 Autopower spectrum per pointing

In this section, we describe tests done with the autopower spectrum (i.e. $j = i$ in 1) for each pointing. The presence of residual systematics that do not correlate out is easier to detect in the autopower spectrum. Thus, we use this test to find the amount and location of contamination expected in the final averaged power. The cylindrical autopower spectrum for each pointing is shown in Fig. 7. As done previously, the pointings are named according to Table 1, and the same declinations are grouped in rows. From the individual panels of Fig. 7, clear systematic contamination is seen inside the observation window i.e. above the horizon line for some pointings. The pointings COSMOS_2, COSMOS_4, COSMOS_0 and COSMOS_3. COSMOS_2, COSMOS_3 and COSMOS_4, which show higher levels of flagging, also show excess power beyond the horizon limit. These results indicate that these pointings are affected more by systematics. We show further evidence of contaminations in Fig. A1, where these pointings have excess power in Stokes V at small k_\perp modes. The first few channels of COSMOS_0 have a high flagging percentage in the subband used (first panel in the third row of Fig. 5). COSMOS_0 also exhibits systematic signatures in the small k_\perp modes beyond the horizon line, as evident from Fig. 7 (and from Fig. A1). The exact reason for the observed contamination is unknown. One possibility is that the observing tracks have the earliest starting times around 15:30 local time (or 13:30 UTC) and also cover local sunset, which affects the data. Fig. 7 does not point to any clear declination dependence of systematics.

As a final test, we show the averaged autopower spectrum of the combined MIGHTEE-COSMOS pointings in Fig. 8. The left panel shows the cylindrical power, while the right panel shows the spherically averaged power spectrum. In the left panel, the excess power seen beyond the horizon limit (black dashed curve) $k_\perp \lesssim 1$ Mpc $^{-1}$ originates from features seen for some of the cases in Fig. 7. The smallest k_\perp -bin (< 0.5 Mpc $^{-1}$) is noise-dominated and




Figure 7. Cylindrical averaged autocorrelation power spectrum of the continuum-subtracted MIGHTEE-COSMOS data. Each panel is the autocorrelation power spectrum of a single pointing, as indicated in the panel name. Each row represents pointings at the same declination. There are visible signatures of systematics present in some pointings, near the blacked dashed curve ($k_{\parallel} \sim 0.3k_{\perp}$, horizon line) for small k_{\perp} values.




Figure 8. **Left:** Cylindrical autopower spectrum from continuum subtracted visibilities for all pointings combined. The dashed and dotted curves are different foreground cuts to avoid systematics affected regions: horizon limit i.e. $k_{\parallel} \sim 0.3k_{\perp}$ (black dashed curve), $k_{\parallel} \sim 0.5k_{\perp}$ (maroon dashed curve), $k_{\parallel} \sim 0.3k_{\perp} + 0.3$ (red dotted curve) and $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ (blue dotted curve). **Right:** Spherically averaged autopower spectrum of continuum-subtracted MIGHTEE-COSMOS data. The black squares show the power spectrum calculated from the data using modes beyond the horizon limit, i.e. $k_{\parallel} \sim 0.3k_{\perp}$, and the red dashed line shows the average thermal noise power from all pointings.

may have subdominant systematics. Insufficient signal-to-noise ratio prevents us from properly determining the major contaminant in this bin. We show the spherical autopower spectrum marked by black squares in the right panel of Fig. 8. We use the modes outside the horizon limit of $k_{\parallel} \sim 0.3k_{\perp}$ (i.e. marked by the black dashed curve on the left panel). The average thermal noise is shown by the red dashed line. The autopower measurements follow the noise except the smallest k -bin, where the dominant contribution comes from the k -modes contaminated by systematics. Removing the noise bias should give us a better estimation of the residual power in the data, which is discussed later.

3.3.2 Effect of changing measurement window

We test whether excluding the contaminated modes beyond the horizon limit significantly improves the 1D power spectrum. We test three different foreground avoidance criteria (marked by different coloured curves in Fig. 8): $k_{\parallel} \sim 0.5k_{\perp}$ (zone beyond the magenta dashed curve), $k_{\parallel} \sim 0.3k_{\perp} + 0.3$ (zone beyond red dotted curve) and $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ (zone beyond the blue dotted curve). The foreground cuts are used to measure the cross-power spectrum (to remove the effect of noise bias) and compare with the horizon limit cuts. Fig. 9 shows the spherical cross-power spectra in different subplots- the top left is the spherical power with the original horizon limit, the top right is $k_{\parallel} \sim 0.3k_{\perp} + 0.3$, bottom left is $k_{\parallel} \sim 0.5k_{\perp}$ and bottom right is $k_{\parallel} \sim 0.5k_{\perp} + 0.5$. It is seen from the bottom right subplot that while there is an improvement in the smallest k bin, the overall power spectrum is still relatively noisy. The other two cases show no marked improvement compared to the horizon cut. The bottom right panel ($\sim 0.5k_{\perp} + 0.5$) does show some improvement in the first bin, but it also excludes a large number of modes. Since more stringent cuts lower the signal-to-noise ratio, with improvements only in the smallest k bin, we use the original measurement window beyond the instrument horizon for our final measurements. It should also be noted from Fig. 8 that most modes beyond the foreground wedge region are clean of systematics and a measurement window, excluding the few contaminated modes at small k_{\perp} , but including




Figure 9. Spherical cross-power spectrum with different foreground cuts shown in the left panel of Fig. 8 – horizon limit of $k_{\parallel} \sim 0.3k_{\perp}$ (black squares, upper left panel), $k_{\parallel} \sim 0.3k_{\perp} + 0.3$ (red stars, upper right panel), $k_{\parallel} \sim 0.5k_{\perp}$ (maroon crosses, lower left panel), and $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ (blue pentagons, lower right panel).

everything beyond the wedge boundary can enhance the signal-to-noise ratio. However, such measurement windows would be specific to this particular data set, and therefore, we use $k_{\parallel} \sim 0.3k_{\perp}$ limit to keep the analysis general.

3.3.3 Power Distribution per k bin

In this section, we examine the statistics of the averaged cross-power spectrum. The Fourier modes are binned into logarithmic k bins, and the resulting distribution in each k mode is plotted in Fig. 10. Two cases are shown with different measurement windows – $k_{\parallel} \sim 0.3k_{\perp}$ (black dashed histograms) and $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ (blue solid histogram). The level of outliers in each k -bin can be detected

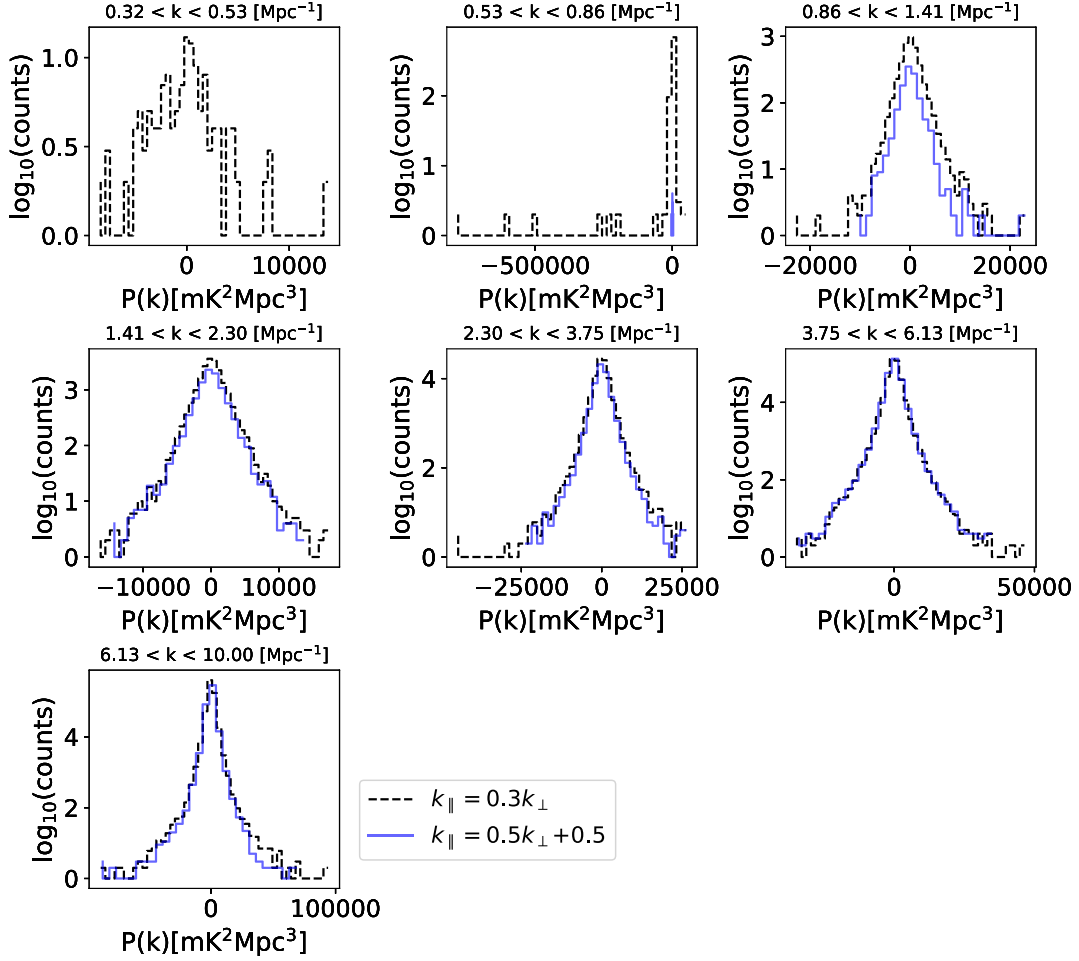


Figure 10. Distribution of the power spectrum per k bin from cross-correlating the odd and even scans for two measurement windows. The black dashed histograms are for the $k_{\parallel} \sim 0.3k_{\perp}$ cut and blue solid histogram is for $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ cut. The latter cut excludes the first k bin, hence it contains only the black curve. The latter cut excludes more modes resulting in difference in the histogram count. Using the more stringent foreground cuts results in an improvement in the smallest k mode, but there is no significant difference seen for the other modes.

from the distribution of power in it. The noise is expected to be Gaussian random, thus heavily skewed distributions indicate the power to be systematics-dominated rather than noise-dominated. The $k_{\parallel} \sim 0.5k_{\perp} + 0.5$ cut does not measure the smallest bin, hence the first panel only contains the black histogram.

The smallest scale that can be measured with both cuts is $\sim 0.7 \text{ Mpc}^{-1}$. The blue histogram in the corresponding bin (second panel in the first row of Fig. 10) has a very small spread. Hence, using the stringent cut improves the measurement in the smallest k -mode by both (also evident from the bottom right panel of Fig. 9). However, the histograms in other bins are both similarly distributed. Thus, the stricter cut does not show a significant improvement over using the horizon limit. Additionally, the smaller measurement window also reduces the number of accessible modes and the counts in each bin. The skewed histogram in the first panel also shows that a large portion of the discernible contamination is from the systematics in the shortest baselines contributing dominantly in this bin. Nevertheless, Fig. 10 reinforces that using a more stringent cut does not improve the measurements considerably. Additionally, the limit $k_{\parallel} \sim 0.3k_{\perp}$ gives a larger number of k modes and sufficiently large counts across them; hence, it is used for our final measurements.

3.4 Dropout tests with different pointings

As evident from Section 3.3.1, some pointings have systematics in the small k_{\perp} modes. Hence, it is plausible that removing these contaminated pointings can improve the overall results, with the obvious caveat that the large amount of data removed will lower the signal-to-noise. Since the MIGHTEE-COSMOS tracks the target continuously for about an hour before the intermediate secondary calibrator scan, time-dependent systematics present in one or more scans may impact the overall results. However, if these contaminations are transient (i.e. present in one scan but absent in the others), they would drop out on correlating different scans. But if present over larger time-scales, the effects may be enhanced (or may not correlate out as effectively) on cross-correlating. Thus, we do drop-out tests by removing each pointing entirely to check for time-dependent systematics in the data.

In Fig. 11, we show the averaged cylindrical power spectra after dropping out two pointings – COSMOS_3 and COSMOS_10. They represent pointings for which Fig. 7 shows clear systematics in the measurement window (upper panel) and a case with no visible contamination (lower panel). We see that combining pointings makes the level of contamination average down, but only very slightly (for instance as seen in the top panel). The ‘cleanest’ results would of




Figure 11. Cylindrical autopower spectrum for dropout tests. Top panel: This shows the result on removal of COSMOS_3, a pointing with clear systematic contamination seen in Fig. 7. While some of the bins show lower levels of excess power, there are still some persistent features due to contaminants in other pointings. Bottom panel: This shows the case on removing COSMOS_10, one of the clean pointings. As expected, there is no clear improvement seen on removing a pointing with no visible contamination. Both cases use continuum-subtracted visibilities.

course include only clean pointings, but given the volume of data that we would lose, it is not feasible for this work.

Fig. 12 shows the 1D cross-power spectrum for the drop-out tests demonstrated in Fig. 11. Similar to Section 3.3.2, we use cross-power to remove the noise bias. We use 7 logarithmic bins for all the cases, but Fig. 12 has the k values slightly displaced for the different samples for better visualization. It is seen that the first k mode (centred at 0.34 Mpc^{-1}) is negatively biased, indicative of some systematics in the smallest k_{\perp} mode. The second bin centred at 0.57 Mpc^{-1} shows improvement on removing COSMOS_3 (though the error bar is still large). For the other bins, the results are sufficiently consistent across all the cases. Investigations into COSMOS_3 show that while most other pointings start observing late afternoon and get only the local sunset, it starts observation near local noon. The observation time may be the cause of the slightly greater amount of contamination in the data. It is also possible that local ionospheric conditions on this particular day were more turbulent. It also shows that subdominant systematics can be a problem with the method used since it is not always obvious over the noise. However, it is expected that with the addition of more data, we can apply more stringent cuts without




Figure 12. Spherically averaged cross-power spectrum for the drop-out tests with individual pointings removed. The samples shown are COSMOS_2 (red diamonds), COSMOS_3 (green crosses), COSMOS_4 (blue circles), COSMOS_10 (navy triangles), COSMOS_6 (magenta right carat), COSMOS_7 (cyan left carat), COSMOS_12 (orange inverted triangles), and COSMOS_14 (maroon stars). The error bars are the 1σ uncertainties for each case.

compromising too much on the signal-to-noise ratio and thus improve the results.

3.5 Effect of continuum subtraction

The visibility domain continuum subtraction described in Section 2.1 could potentially cause H I signal loss by removing H I along with the continuum emission. In Fig. 8, the amplitude of the foreground wedge is an order of magnitude lower than expected. The continuum emission from foregrounds, the dominant contributor to the wedge power, is removed from the visibilities used here for H I IM.

The visibility domain subtraction uses the model created from the deep MIGHTEE continuum images and subtracts the continuum model from self-calibrated data. The obtained residuals are used for our analysis, so we calculate the ratio of this continuum-subtracted visibility with the visibility containing the continuum (or foregrounds). If the process does not remove any significant amount of H I, the ratio in the window should be ~ 1 , while a significantly smaller value will be seen in the wedge. We also measured the cylindrical power spectrum of the subtracted continuum model. In the absence of H I in the model, we would expect negligible power in the region outside the wedge.

The left panel of Fig. 13 shows the ratio between the cylindrical power spectra of the continuum-subtracted visibilities and that of visibilities pre-subtraction. Any significant signal removal due to subtraction would result in the ratio being $\ll 1$ in the H I measurement window. However, the ratio is almost consistently 1 in the region beyond the white dashed line, with the smallest value in the foreground wedge, where most of the removed power resides. The right panel of Fig. 13 shows the power spectrum of the subtracted continuum model. We see here that inside the foreground wedge, the power spectrum has a value $\gtrsim 10^5 \text{ mK}^2 \text{ Mpc}^3$. Beyond the horizon limit (beyond the white dashed line), the values are very small ($< 10^{-2} \text{ mK}^2 \text{ Mpc}^3$), a much smaller value than that expected for H I signal.

Continuum subtraction removes a sky model dominated by the spectrally smooth foregrounds but preserves the spectral behaviour




Figure 13. Left: Ratio of cylindrical autopower spectra of continuum subtracted visibilities and visibilities presubtraction, averaged for all pointings. The ratio in the foreground wedge is < 1 , indicating a significant removal of continuum power. The region beyond the white dashed line shows a value of ~ 1 , indicating the measurement window is unaffected by continuum subtraction. Right: Cylindrical power of the continuum model. The power in the measurement window (beyond the white dashed region) is $< 10^{-2} \text{ mK}^2 \text{ Mpc}^3$, indicating negligible power present in the model outside the wedge.

of H I sources. This is essential for both IM and galaxy science. The method has already been tested for MIGHTEE data used for H I studies (for example in Sinigaglia et al. 2022; Heywood et al. 2024). Fig. 13 shows that the emission removed is almost exclusively confined to the wedge, which is evidence that the sky model used for the visibility domain continuum subtraction has well-modelled emission from the wedge only. It works well in reducing the foreground emission, without significantly impacting the underlying H I.

From our tests, we could not find any evidence for H I signal loss through the continuum removal process. Generally, the full extent of H I signal loss can only be fully quantified through end-to-end simulations, which is beyond the scope of this paper but should be performed for future analysis.

4 RESULTS

In this section, we discuss the results obtained from the MIGHTEE-COSMOS data using the subband centred at 986 MHz ($z \sim 0.44$). We average all 15 pointings in the power spectrum domain and expect lower signal-to-noise ratio than a single pointing with the same integration time. Nevertheless, owing to the excellent data quality and high-frequency resolution, this data is used to set the first upper limits on the H I power spectrum using the MIGHTEE survey. We choose the frequency band centred at $z \sim 0.44$ and the scales are between ~ 0.5 and $\sim 10 \text{ Mpc}^{-1}$ to avoid regions with very high noise.

We present the cross-correlation and autocorrelation power spectra obtained from the gridded visibilities. The latter acts as a conservative upper limit, as it is more likely to contain additive observational systematics. The cross-correlation amplitude is independent of the uncorrelated noise bias. However, negative systematics could artificially suppress the amplitude of the power spectrum, as outlined in (Morales, Pober & Hazelton 2023). Comparing auto and cross-correlations, we verify the impact of noise and systematics and report consistent upper limits on the H I power spectrum within the limited SNR of the MIGHTEE data.

The power is calculated separately for each pointing and incoherently averaged using equation (2). Fig. 14 shows the cylindrical cross-power spectrum for the combined MIGHTEE-COSMOS data centred at $z \sim 0.44$. Owing to the fine frequency resolution of




Figure 14. Incoherently averaged cylindrical cross-power spectrum using continuum subtracted data for all pointings. The cross-power per pointing is the cross-correlation of alternate target scans to remove noise bias, which are then averaged. The black dashed curve shows the horizon limit of $k_{\parallel} \sim 0.3k_{\perp}$, which gives the measurement window for the 1D power.

MIGHTEE, we can probe higher k_{\parallel} modes compared to DEEP2 (Paul et al. 2023). It is seen that the measurement window beyond the horizon limit (black dashed curve) does not show any dominant systematic contribution. The cross-correlation also removes some of the systematic contributions seen at the lower k_{\perp} modes in Fig. 8. The foreground wedge region is less pronounced due to continuum subtraction.

We use the foreground avoidance method for calculating the 1D power spectrum averaging data in the region above the horizon limit $k_{\parallel} \sim 0.3k_{\perp}$ (black dashed line in Fig. 14). We compute the 3D power spectra in 7 logarithmic k bins. Owing to the the smallest k_{\perp} bin being highly noisy, we also exclude $k_{\perp} \leq 0.5 \text{ Mpc}^{-1}$. The final incoherently averaged 1D power spectra at $z \sim 0.44$ for the MIGHTEE-COSMOS observation is shown in Fig. 15. The indigo squares indicate the cross-power spectrum, and blue dashed line marks the amplitude




Figure 15. Incoherently averaged 1D power spectrum from the MIGHTEE-COSMOS data at $z \sim 0.44$. The measurements use modes beyond the horizon limit, i.e. the black dashed line in Fig. 14. The indigo squares show the cross-power spectrum, and the blue dashed curve shows the autopower spectrum. The autopower spectrum is set as the absolute upper limit on the HI power spectrum obtained in this work. The values for the HI power obtained in (Paul et al. 2023) at $z \sim 0.44$ are shown in grey circles for comparison.

Table 2. HI power spectrum constraints at $z \sim 0.44$ obtained from MIGHTEE-COSMOS data after incoherent averaging all pointings. The measurement and errors are tabulated for the cross-power spectrum, while the upper limits indicate the values obtained from the bias-subtracted autopower spectrum. All values are for the same k bins.

k (Mpc^{-1})	$P(k)$ ($\text{mk}^2 \text{Mpc}^3$)	σ_P	$P(k)/\sigma_P$ ($\text{mk}^2 \text{Mpc}^3$)	Upper limit ($\text{mk}^2 \text{Mpc}^3$)
0.69	330.14	122.46	2.69	764.50
0.93	95.72	103.44	0.93	130.21
1.35	18.29	16.49	1.11	59.67
1.96	19.79	5.00	3.96	25.82
2.85	3.15	1.80	1.75	7.29
4.16	0.22	0.69	0.32	4.30
5.99	0.47	0.29	1.63	4.83

of the noise-bias removed autopower spectrum. We also show the spherical power spectrum amplitude at $z \sim 0.44$ reported in Paul et al. (2023) by grey circles. The measured values of the power spectra are also shown in Table 2, between $0.5 \text{ Mpc}^{-1} \lesssim k \lesssim 10 \text{ Mpc}^{-1}$.

We find that the values of the auto and cross-power are consistent within 2σ , except for the last two k bins, which are generally noise dominated, as can be seen in the right panel of Fig. 8. We think the discrepancy is caused by noise and residual systematics in the autopower spectrum in these bins, rather than a suppression in the amplitude of the cross-power. The consistency of the results demonstrates the robustness of our measurements, since the auto and cross-correlations are sensitive to different systematic and noise effects.

The measured cross-power spectrum shows a drop with increasing k , which is expected from the combination of shot noise and attenuation by the Finger-of-God (FoG) at high k_{\parallel} . Compared to the detection obtained in Paul et al. (2023), the power measured at a similar range of k values is higher in this work. This is due to the lower SNR of the data as our analysis combines 15 overlapping pointings

via incoherently averaging, whereas Paul et al. (2023) analysed one deep pointing.

5 SUMMARY

In this work, we demonstrate the capability for measuring the HI IM signal on quasi-linear scales from the visibilities of DR-1 MIGHTEE spectral line observations (Heywood et al. 2024). We use a small subband centred at redshift $z \sim 0.44$ covering around 4 square degrees of the COSMOS field. COSMOS is a well-known extragalactic deep field with a rich population of interesting radio sources for continuum science. However, in our context, these sources contribute to the foreground emission which could potentially contaminate the HI signal. In this work, we found that these emissions are well confined within the horizon scales of the wedge due to the high calibration accuracies, thus validating our approach of foreground avoidance in the power spectrum estimation. We compute both the autocorrelation and the cross-correlation powers from interleaving scans for individual pointings to identify noise and systematic effects. Using cross-correlation, we derive the upper limit of $29.8 \text{ mK}^2 \text{ Mpc}^3$ from the 2σ value which is consistent with the autocorrelation value of $25.82 \text{ mK}^2 \text{ Mpc}^3$ on $k \sim 2 \text{ Mpc}^{-1}$ on the HI power spectrum. The results are within two orders of magnitude of Paul et al. (2023) at the same redshift, even at the smallest k values (despite significantly higher noise equivalent).

A few important things need to be considered in this context:

(i) The data spans a large area over multiple pointings centred at slightly different right ascensions and declinations. Thus, the averaging is done in the power spectrum domain rather than in the visibility domain. Visibility averaging often makes residual errors easier to localize and remove, which cannot be done in this case.

(ii) The observations for MIGHTEE-COSMOS start late afternoon to late evening, none start late at night. Radio surveys are mostly done late at night to avoid ambient turbulence caused by the sun. However, inspection of the data showed that while some contamination is evident, the overall quality is not adversely affected.

(iii) Direction-dependent effects have not been corrected for in this data set. We do not expect this to be an issue for the tightly packed mosaiced pointings used for the L-band MIGHTEE-COSMOS observations. However, we will perform these corrections when combining data for other MIGHTEE fields with wider pointing spreads.

(iv) To assess the quality and usability of the data for IM, several tests were done as discussed in Section 3. The main conclusion from these tests is that residual systematics can vary considerably between pointings (and even from one scan to the other in the same pointing). However, on combination and cross-correlation, most systematic contributions reduce considerably (Fig. 14).

(v) We select the k modes conservatively, excluding those potentially containing systematics obscured by very high noise.

(vi) The limiting issue of identifying systematics due to lower signal-to-noise ratio for this data set will be addressed in future works incorporating more data.

ACKNOWLEDGEMENTS

The authors acknowledge Ian Heywood for the initial data calibration. The authors also thank the anonymous referees for their insightful comments that helped improve the quality of this paper. AM is supported by a UK Research and Innovation Future Leaders Fellowship grant [MR/V026437/1]. LW is a UK Research and Innovation Future Leaders Fellow [MR/V026437/1]. MGS

acknowledges support from the South African Radio Astronomy Observatory and National Research Foundation (grant 84156). MJJ acknowledge generous support from the Hintze Family Charitable Foundation through the Oxford Hintze Centre for Astrophysical Surveys, support of the STFC consolidated grant [ST/S000488/1] and [ST/W000903/1] and from a UKRI Frontiers Research Grant [EP/X026639/1]. 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. We acknowledge the use of the Ilifu cloud computing facility (<https://www.ilifu.ac.za/>), a partnership between the University of Cape Town, the University of the Western Cape, Stellenbosch University, Sol Plaatje University and the Cape Peninsula University of Technology. The Ilifu facility is supported by contributions from the Inter-University Institute for Data Intensive Astronomy (IDIA – a partnership between the University of Cape Town, the University of Pretoria and the University of the Western Cape), the Computational Biology division at UCT and the Data Intensive Research Initiative of South Africa (DIRISA). The authors acknowledge the Centre for High Performance Computing (CHPC), South Africa, for providing computational resources to this research project. The authors thank the reviewer and the scientific editor for their insightful comments that have helped in improving this work. AM also thanks Sumanjit Chakraborty, Steve Cunnington, Keith Grainge and Amadeus Wild for helpful discussions.

DATA AVAILABILITY

The visibility data are available from the SARAO archive (https://keycloak.sarao.ac.za/auth/realms/SKASA/login-actions/authenticate?client_id=archive.sarao.ac.za&tab_id=_BjAC2k2F90) by searching for the capture block IDs listed in Table 1.

Softwares: This work relies on the Python programming language (<https://www.python.org/>). The packages used here are ASTROPY (<https://www.astropy.org/>; ASTROPY Collaboration 2013; Price-Whelan et al. 2018), NUMPY (<https://numpy.org/>), SCIPY (<https://www.scipy.org/>), MATPLOTLIB (<https://matplotlib.org/>), Common Astronomy Software Applications (CASA) (https://casaguides.nrao.edu/index.php?title=Main_Page). Some analyses were done using HIIMTOOL (<https://github.com/zhaotingchen/hiimtool>).

REFERENCES

- Abbott T. M. C. et al., 2022, *Phys. Rev. D*, 105, 023520
- Abdurashidova Z. et al., 2022, *ApJ*, 925, 221
- Adams N. J., Bowler R. A. A., Jarvis M. J., Häußler B., McLure R. J., Bunker A., Dunlop J. S., Verma A., 2020, *MNRAS*, 494, 1771
- Adams N. J., Bowler R. A. A., Jarvis M. J., Häußler B., Lagos C. D. P., 2021, *MNRAS*, 506, 4933
- Alam S. et al., 2021, *Phys. Rev. D*, 103, 083533
- Amiri M. et al., 2023, *ApJ*, 947, 16
- Amiri M. et al. 2024, *The Astrophysical Journal* 923 preprint(arXiv:2309.04404)
- Anderson C. J. et al., 2018, *MNRAS*, 476, 3382
- Astropy Collaboration, 2013, *A&A*, 558, A33
- Bandura K. et al., 2014, in Stepp L. M., Gilmozzi R., Hall H. J., eds, Proc. SPIE Conf. Ser. Vol. 9145, Ground-based and Airborne Telescopes V. SPIE, Bellingham, p. 914522
- Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, *MNRAS*, 461, 3135
- Battye R. A., Davies R. D., Weller J., 2004, *MNRAS*, 355, 1339
- Bharadwaj S., Nath B. B., Sethi S. K., 2001, *JA&A*, 22, 21
- Bowler R. A. A., Jarvis M. J., Dunlop J. S., McLure R. J., McLeod D. J., Adams N. J., Milvang-Jensen B., McCracken H. J., 2020, *MNRAS*, 493, 2059
- Chakraborty A. et al., 2021, *ApJ*, 907, L7
- Chan K. C. et al., 2022, *Phys. Rev. D*, 106, 123502
- 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
- Chapman E. et al., 2015, in Bourke T.L. et al., eds, *Advancing Astrophysics with the Square Kilometre Array (ASKA14)* SISSAp. 5
- Chen Z., Wolz L., Spinelli M., Murray S. G., 2021, *MNRAS*, 502, 5259
- Chen Z., Wolz L., Battye R., 2022, *MNRAS*, 518, 2971
- Chen Z., Chapman E., Wolz L., Mazumder A., 2023, *MNRAS*, 524, 3724
- Cleary K. A. et al., 2022, *ApJ*, 933, 182
- Cunnington S., Wolz L., Pourtsidou A., Bacon D., 2019, *MNRAS*, 488, 5452
- Cunnington S., Irfan M. O., Carucci I. P., Pourtsidou A., Bobin J., 2021, *MNRAS*, 504, 208
- Cunnington S. et al., 2022, *MNRAS*, 518, 6262
- DESI Collaboration, 2022, *AJ*, 164, 207
- Datta A., Bowman J. D., Carilli C. L., 2010, *ApJ*, 724, 526
- Dawson K. S. et al., 2016, *AJ*, 151, 44
- de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, *Exp. Astron.*, 35, 25
- Drinkwater M. J. et al., 2010, *MNRAS*, 401, 1429
- Driver S. P. et al., 2011, *MNRAS*, 413, 971
- Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2017, *MNRAS*, 470, 1849
- Furlanetto S. R., Peng Oh S., Briggs F. H., 2006, *Phys. Rep.*, 433, 181,
- Goedhart S., 2020a, MeerKAT specifications. <https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/277315585/MeerKAT+specifications> (accessed 25 June 2025).
- Goedhart S., 2020b, Radio Frequency Interference (RFI), <https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/305332225/Radio+Frequency+Interference+RFI> (accessed 25 May 2025).
- Hale C. L. et al., 2025, MIGHTEE: the continuum survey Data Release 1 *MNRAS*, 536, p. 2211
- Heywood I. et al., 2021, *MNRAS*, 509, 2150
- Heywood I. et al., 2024, *MNRAS*, 534, 76
- Hugo B. V., Perkins S., Merry B., Mauch T., Smirnov O. M., 2022, in Ruiz J. E., Pierfederici F., Teuben P., eds, ASP Conf. Ser. Vol. 532. Astronomical Data Analysis Software and Systems XXX ASP Conf. Ser., San Francisco, p. 541
- Jarvis M. et al., 2016, in Taylor R., Camilo F., Leeuw L., Moodley K., eds, Proceedings of MeerKAT Science: On the Pathway to the SKA. Proceedings of Science SISSA. p. 006
- Jelić V. et al., 2008, *MNRAS*, 389, 1319
- Jonas J., 2018, in Taylor R., Camilo F., Leeuw L., Moodley K., eds, Proceedings of MeerKAT Science: On the Pathway to the SKA Proceedings of Science SISSA. p. 001
- Kovetz E. D. et al., 2017, preprint (arXiv:1709.09066)
- Laporte N. et al., 2017, *ApJ*, 837, L21
- Liu A., Parsons A. R., Trott C. M., 2014, *Phys. Rev. D*, 90, 023018
- Liu A., Zhang Y., Parsons A. R., 2016, *ApJ*, 833, 242
- Masui K. W. et al., 2013, *ApJ*, 763, L20
- Matshawule S. D., Spinelli M., Santos M. G., Ngobese S., 2021, *MNRAS*, 506, 5075
- Mazumder A., Datta A., Chakraborty A., Majumdar S., 2022, *MNRAS*, 515, 4020
- Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
- Morales M. F., Hewitt J., 2004, *ApJ*, 615, 7
- Morales M. F., Wyithe J. S. B., 2010, *ARA&A*, 48, 127
- Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, *ApJ*, 752, 137
- Morales M. F., Beardsley A., Pober J., Barry N., Hazelton B., Jacobs D., Sullivan I., 2018, *MNRAS*, 483, 2207
- Morales M. F., Pober J., Hazelton B. J., 2023, *MNRAS*, 525, 2834
- Offringa A. R., de Bruyn A. G., Biehl M., Zaroubi S., Bernardi G., Pandey V. N., 2010, *MNRAS*, 405, 155
- Paciga G. et al., 2011, *MNRAS*, 413, 1174
- Parsons A., Pober J., McQuinn M., Jacobs D., Aguirre J., 2012, *ApJ*, 753, 81

- Parsons A. R. et al., 2014, *ApJ*, 788, 106
- Paul S., Santos M. G., Townsend J., Jarvis M. J., Maddox N., Collier J. D., Frank B. S., Taylor R., 2021, *MNRAS*, 505, 2039
- Paul S., Santos M. G., Chen Z., Wolz L., 2023, preprint (arXiv:2301.11943)
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Pourtsidou A., Bacon D., Crittenden R., 2017, *MNRAS*, 470, 4251
- Price-Whelan A. M. et al., 2018, *AJ*, 156, 123
- Ratcliffe S., 2020, SDP Pipelines Overview, <https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/338723406/SDP+pipelines+overview> (accessed 25 May 2025).
- SKA Cosmology SWG Bacon D. J. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e007
- Santos M. et al., 2016, in Taylor R., Camilo F., Leeuw L., Moodley K., eds, *MeerKAT2016 MeerKAT Science: On the Pathway to the SKA*. p. 32 preprint (arXiv:1709.06099)
- Sinigaglia F. et al., 2022, *ApJ*, 935, L13
- Spinelli M., Zoldan A., De Lucia G., Xie L., Viel M., 2020, *MNRAS*, 493, 5434
- Spinelli M., Carucci I. P., Cunnington S., Harper S. E., Irfan M. O., Fonseca J., Pourtsidou A., Wolz L., 2021, *MNRAS*, 509, 2048
- Switzer E. R. et al., 2013, *MNRAS*, 434, L46
- Switzer E. R., Anderson C. J., Pullen A. R., Yang S., 2019, *ApJ*, 872, 82
- The Dark Energy Survey Collaboration, 2005, The Dark Energy Survey. preprint(astro-ph/0510346)
- Trott C. M., Wayth R. B., 2016, *Publ. Astron. Soc. Aust.*, 33, e019
- Trott C. M., Wayth R. B., Tingay S. J., 2012, *ApJ*, 757, 101
- Trott C. M. et al., 2020, *MNRAS*, 493, 4711
- Vedantham H., Shankar N. U., Subrahmanyam R., 2012, *ApJ*, 745, 176
- Visbal E., Trac H., Loeb A., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 010
- Wang J. et al., 2021, *MNRAS*, 505, 3698
- Whittam I. H. et al., 2023, *MNRAS*, 527, 3231
- Wolz L., Abdalla F. B., Blake C., Shaw J. R., Chapman E., Rawlings S., 2014, *MNRAS*, 441, 3271
- Wolz L. et al., 2021, *MNRAS*, 510, 3495
- Wyithe J. S. B., Loeb A., Geil P. M., 2008, *MNRAS*, 383, 1195
- York D. G. et al., 2000, *AJ*, 120, 1579
- Yue B., Ferrara A., Pallottini A., Gallerani S., Vallini L., 2015, *MNRAS*, 450, 3829

APPENDIX A: EXCESS POWER IN STOKES V

Since the circular polarization of the sky is assumed minimal, the Stokes V power spectrum is a good approximation to the noise power for each pointing. The ratio of the simulated thermal noise power to the Stokes V cylindrical power spectrum per pointing is shown in Fig. A1.

There is a structure spread across the lowest part of the k_{\perp} - k_{\parallel} plane, similar to the foreground wedge found in Stokes I. The ratio of the two powers is ~ 1 outside this region, making the Stokes V and noise powers consistent. The excess is an indication of power leakage into Stokes V for MIGHTEE COSMOS data. It is seen from Fig. A1 that the feature is present across all pointings, showing maximum excess in Stokes V power in the lowest part of k_{\perp} - k_{\parallel} plane (i.e. low delays and very short baselines). The enhanced noise and use of continuum-subtracted data make the feature disappear at higher k_{\perp} . The leakage appears in modes excluded in the power spectrum estimation using foreground avoidance. The modes being used in our final result are reasonably noise-like. Thus, this systematic feature does not affect our final results. Detailed studies into the causes and extent of the leakage for the full MIGHTEE data and its effects on including contaminated modes are deferred to future works.

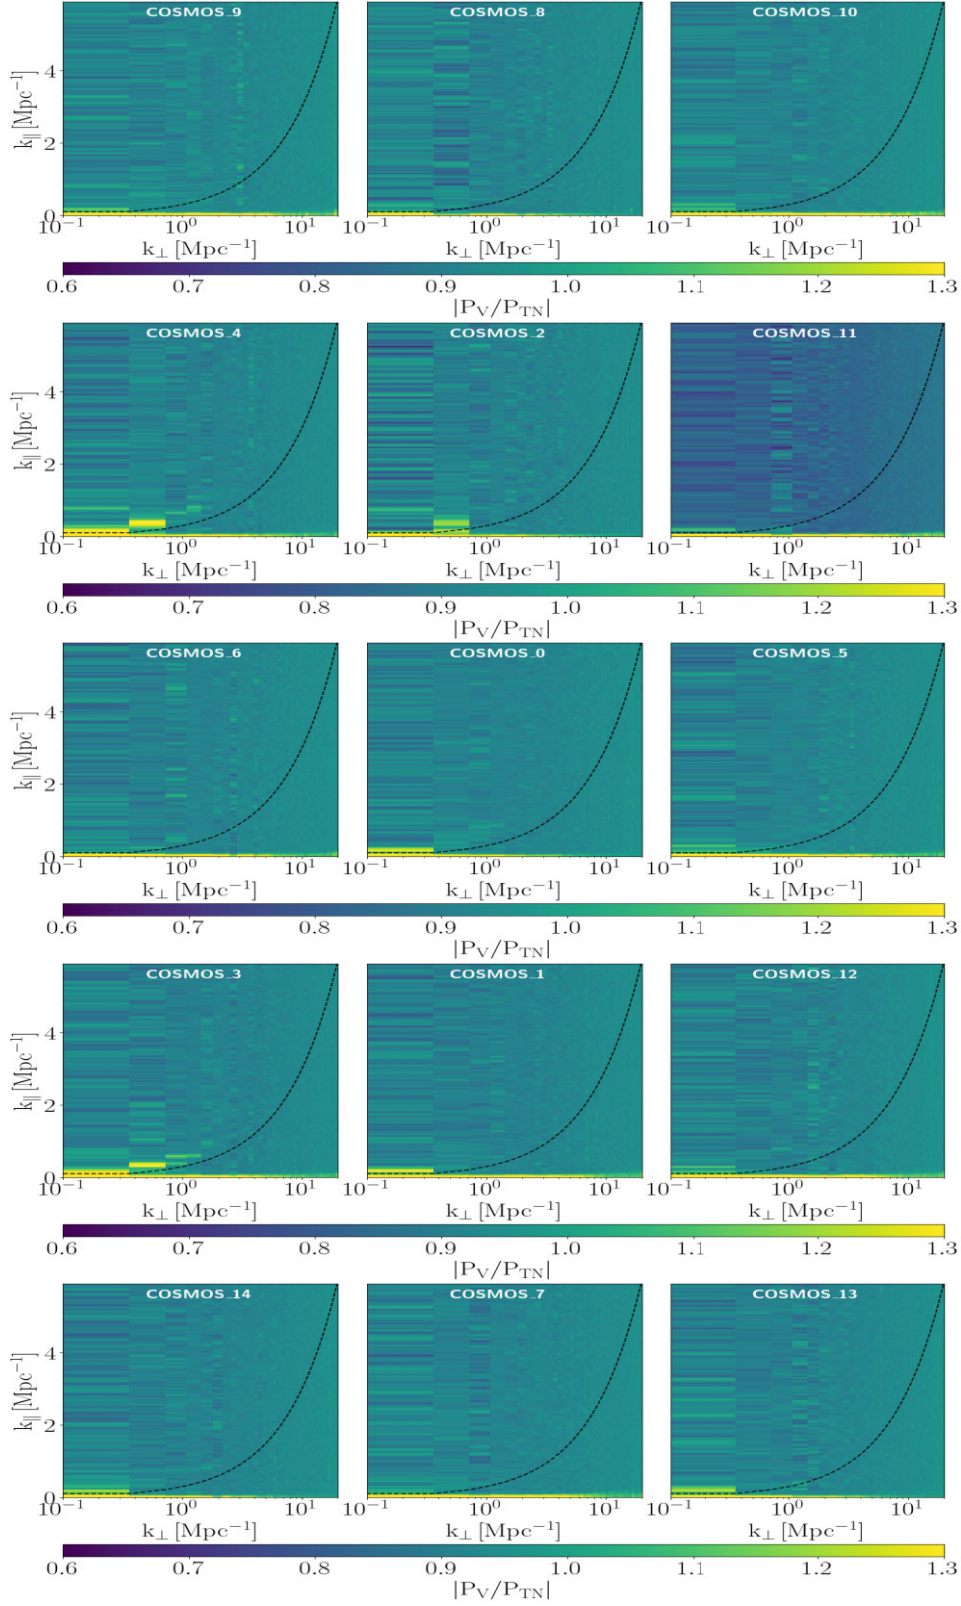


Figure A1. Ratio of the cylindrical power spectrum Stokes V to simulated thermal noise power spectrum for individual pointings. An excess power in Stokes V is seen, particularly at the smallest k_{\perp} modes, similar to the foreground wedge in Stokes I.

This paper has been typeset from a \LaTeX file prepared by the author.